Congestion Control in Local Area Networks

F. Tobagi, C. Fraleigh, W. Noureddine

Outline

- Background
- Approach
  - Simulation model
  - Traffic models
- Experiments
  - Illustrate the need for congestion control (802.3x)
  - Illustrate the need for congestion control based on class of service (CoS)
  - Illustrate the need for congestion control based on destination address
- Conclusions

Background

Introduction

- Extended LANs
  - large scale
  - hundreds of users
  - mix of technologies (e.g. 10, 100, 1000 Mbps segments)
  - supports multiple classes of service (possibly with low delay requirements)
  - congestion may be a problem

Effects of Congestion

- Packet buildup in buffers leading to...
  - increased delay
  - packet loss
    - inefficient use of network resources
      - bandwidth
      - buffer space
      - processing power
    - need for retransmissions

Congestion in LANs

- Congestion in LANs is short-term in nature
  - Generally, LANs' capacity is over-provisioned
  - Long term congestion is dealt with by higher layers (e.g., TCP, etc...)

Sources of LAN Congestion

- Burstiness in traffic
  - demand temporarily exceeds available resources at some point in the network e.g.,
    - Traffic Merging

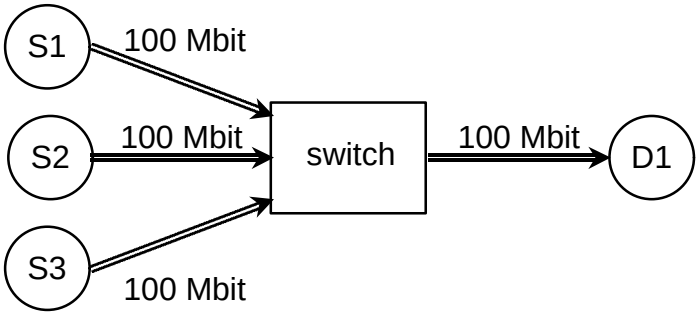

    - Rate Mismatch

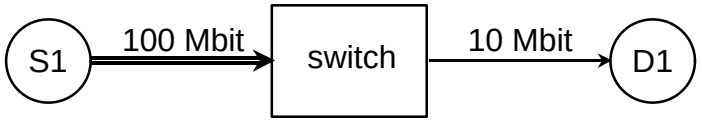

Congestion Control Mechanism (1)

- One may define a congestion mechanism in terms of a minimum of three components (steps)
  - Congestion detection, e.g.,
    - based on buffer occupancy
    - based on the rate at which buffer is filling up
  - Notification
    - which switch to notify, e.g.,
      - all neighboring switches
      - switches that are currently sending packets to the congested buffer
    - what information, e.g.,
      - class of service of congested buffer
      - MAC address information

Congestion Control Mechanism (2)

- Response to notification, e.g.,
  - block/unblock (e.g. IEEE802.3x)
  - rate control
- May extend above functionality to end stations

Congestion Control Mechanism

Congestion Detection

- Performed at switch output buffers
  - High Threshold
    - Congestion is considered to have occurred when buffer occupancy exceeds the high threshold
    - Needs to be low enough to handle packets that arrive before congestion control actions take effect
  - Low Threshold
    - Congestion is considered to be relieved when buffer occupancy falls below the low threshold
    - Needs to be high enough to prevent starvation before congestion control actions are reversed

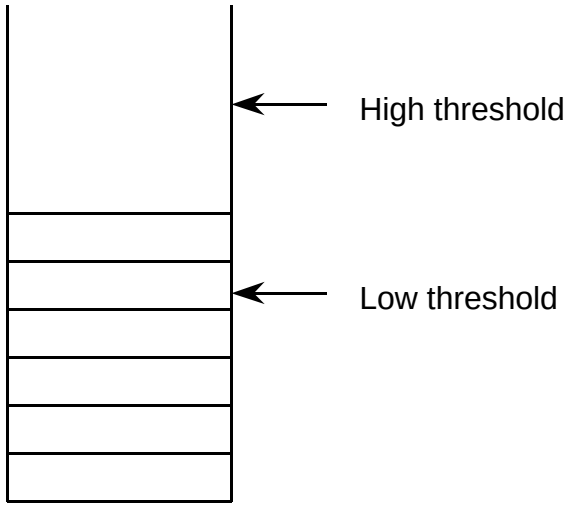

Congestion Control Mechanism

Notification Information

- No specific information
  - block/unblock all traffic IEEE 802.3x
- Class of Service information
  - block/unblock specified priority class
    - the class of service of the congested buffer is readily available
- Destination address information
  - block/unblock traffic to specified destination addresses
    - information about all MAC addresses that are reached through the congested port is available in the filtering database
    - can also look at packets in the congested buffer and extract destination addresses
- Similar notification messages can be sent asking for rate control instead of blocking

Approach

Simulation model

Traffic models

Approach

Simulation Model

Simulation Model

- Simulator for switched Ethernet LAN
  - Uses full-duplex links
  - Supports 10, 100 and 1000 Mbps links
  - Supports multiple traffic classes
  - Switch model
    - non-blocking
    - implements output buffering
    - uses a separate queue for each class of service
      - service discipline is *highest priority class first*
    - will be extended to handle different switch models

Approach

Traffic Models

Uniform-Fixed

Uniform-Uniform

Self-Similar

Video

Uniform-Fixed Data Traffic Model

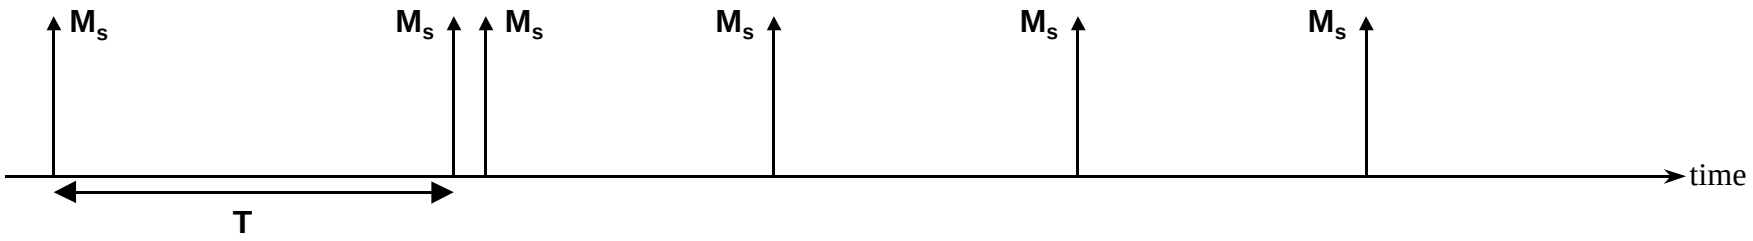

- Uniformly distributed arrival times
  - between 0 and 2T
- Fixed burst size M_s (bytes)
  - Range for M_s = 6,000 … 96,000
- Load G_s (bits per second) = $8\ M_s / T$

Uniform-Uniform Data Traffic Model

- Uniformly distributed arrival times
  - between 0 and 2T
- Random burst size
  - Uniformly distributed between 64 and 2 M_s
    - $X_i \sim U(64, 2M_s)$
    - $\overline{X} = M_s - 32$
- Load G_s (bits per second) = $8(M_s - 32)/ T \cong 8 M_s /T$

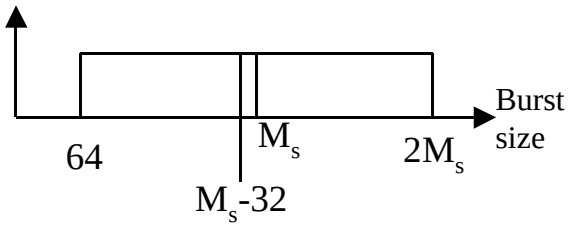

Self-Similar Data Traffic Model

- Accurately models real backbone Ethernet traffic
- May be artificially generated by the aggregation of many (100 or more) bursty data sources
  - X and T have the Pareto distribution (characterized by a heavy tail - with very large variance!)

Video Traffic Model

- Star Trek video trace
- MPEG1, 1.5 Mbit/sec
- VBR

Experiments

A. Illustrate the need for congestion control

B. Illustrate the need for congestion control based on CoS

C. Illustrate the need for congestion control based on destination address

Experiments

Illustrate the need for congestion control

A.1 Traffic Merging

A.2 Rate Mismatch

Traffic Merging Topology

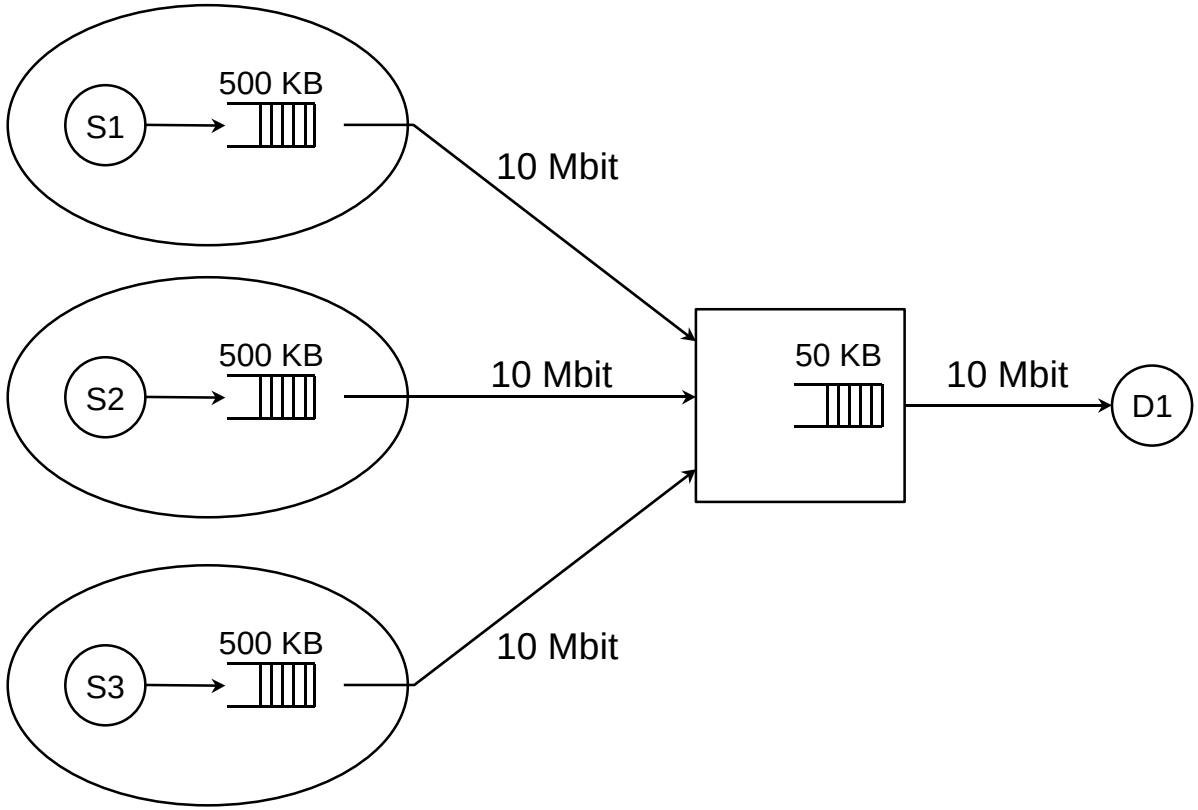

Traffic Merging (1)

- Traffic
  - Uniform-uniform traffic
    - Burst size range 6,000…48,000 bytes
      - same value is used for all 3 sources
    - Data rate range 1 Mbps to 3 Mbps per source
      - same value is used for all 3 sources
  - Self-similar traffic
    - Burst size range 6,000…48,000 bytes
      - same value is used for all 3 sources
    - Data rate range 1 Mbps to 3 Mbps per source
      - same value is used for all 3 sources

Traffic Merging (2)

- Congestion control mechanism
  - watermark-based congestion detection
    - low threshold = 70%
    - high threshold = 80%
  - notification information
    - block/unblock with no specific information
- Measures
  - packet loss rate

Traffic Merging

Packet Loss without Congestion Control

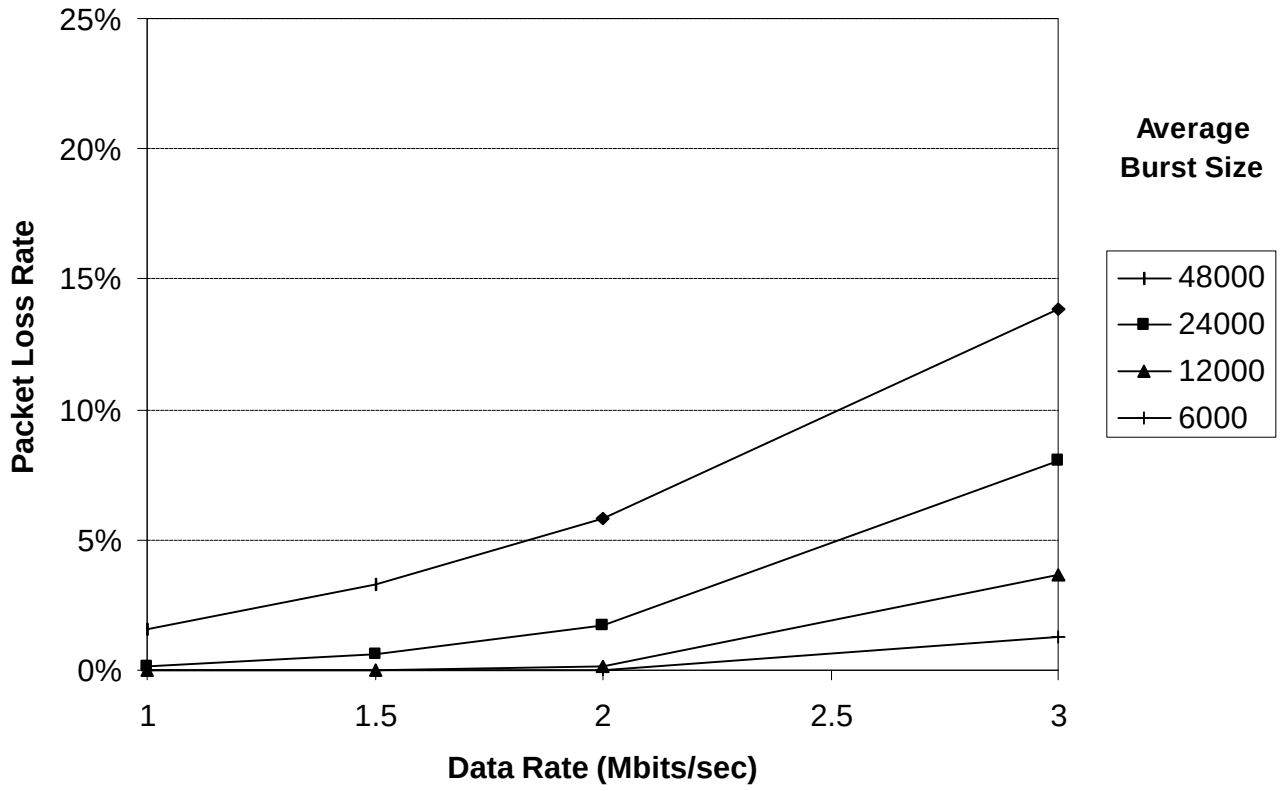

- Uniform-uniform traffic model
- Data rate is shown per source
- 50 KB buffer becomes congested and drops packets

Traffic Merging

Packet Loss using XON/XOFF

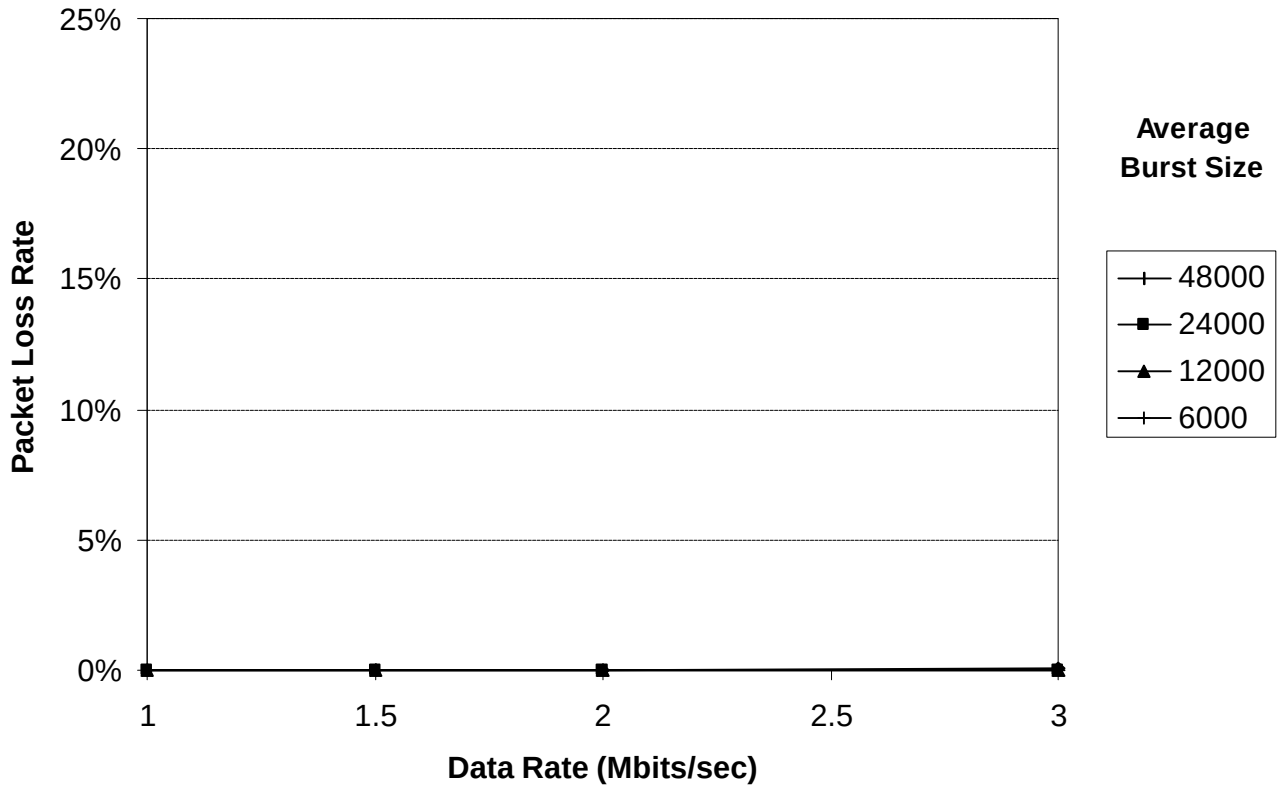

- Uniform-uniform traffic model
- Data rate is shown per source
- High threshold 80%
- Low threshold 70%

Traffic Merging

Packet Loss without Congestion Control

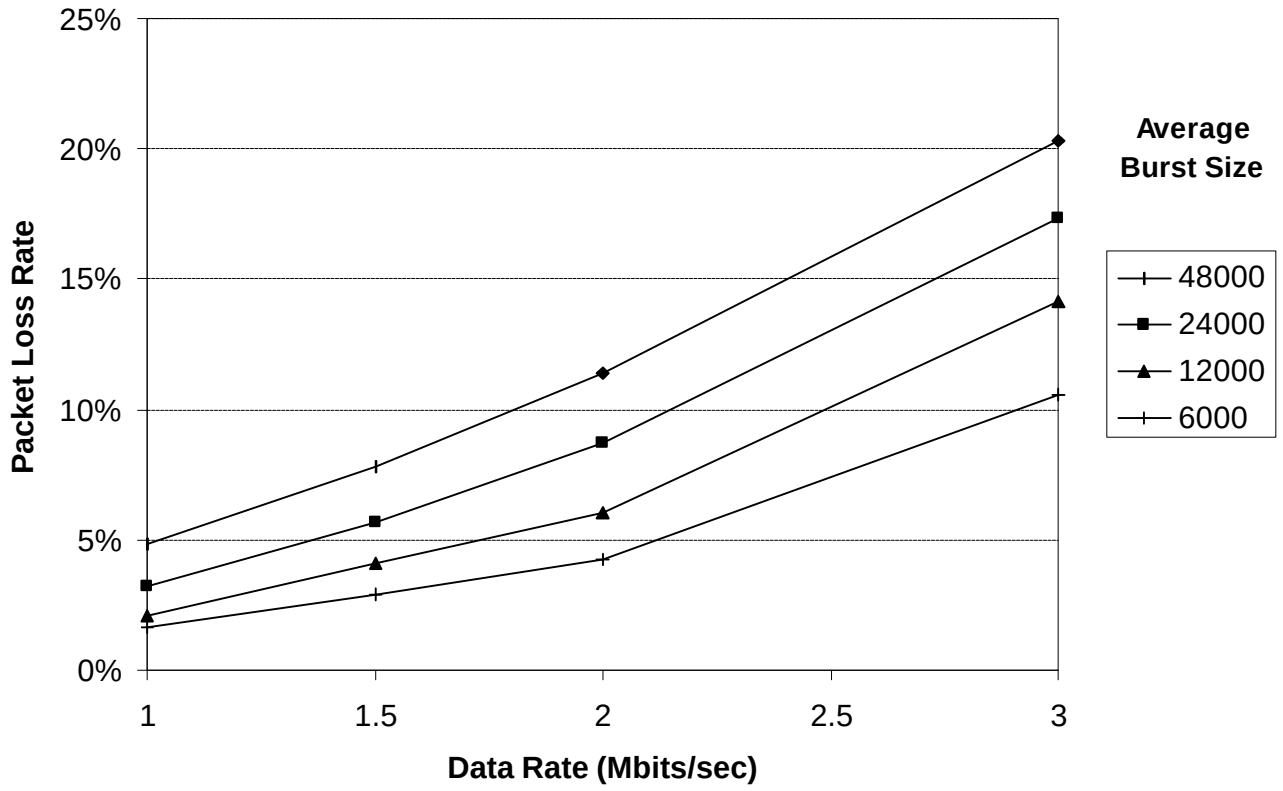

- Self-similar traffic model
- Data rate is shown per source
- 50 KB buffer becomes congested and drops packets

Traffic Merging

Packet Loss using XON/XOFF

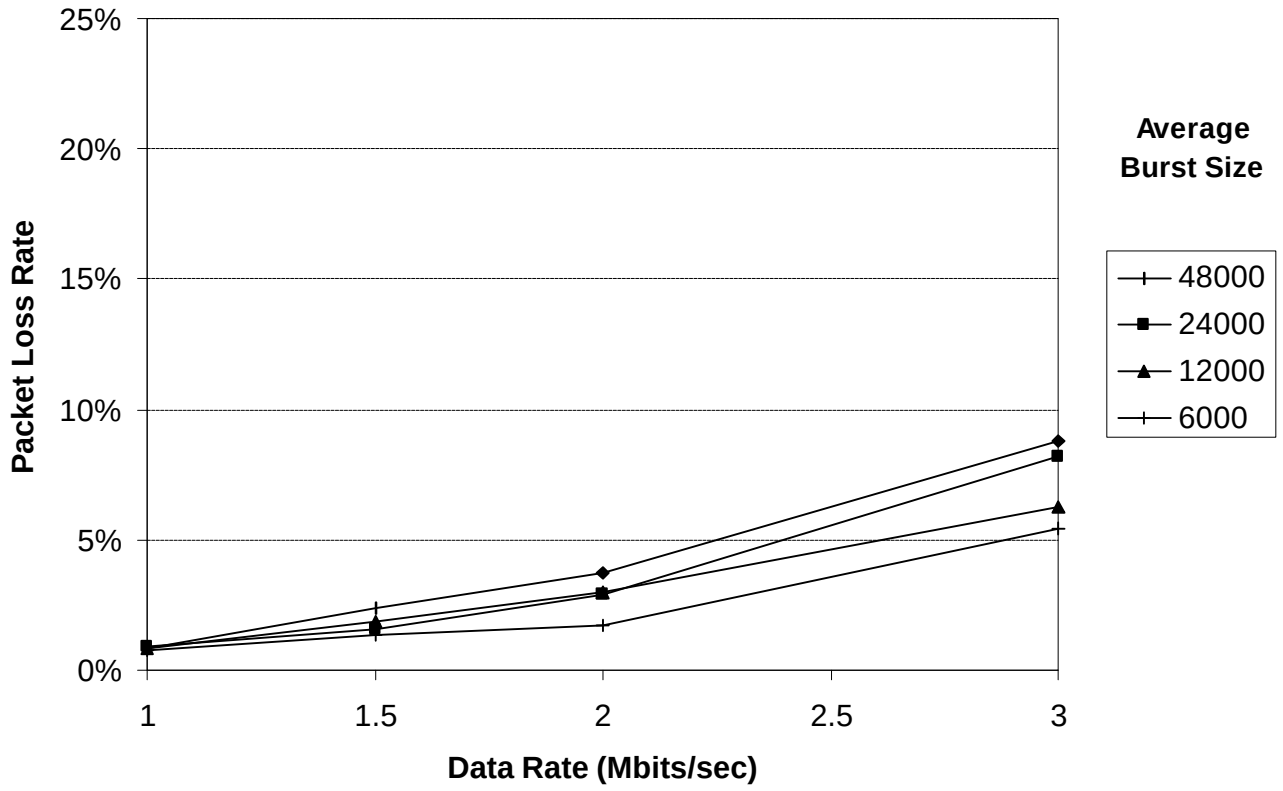

- Self-similar traffic model
- Data rate is shown per source
- High threshold 80%
- Low threshold 70%

Rate Mismatch Topology

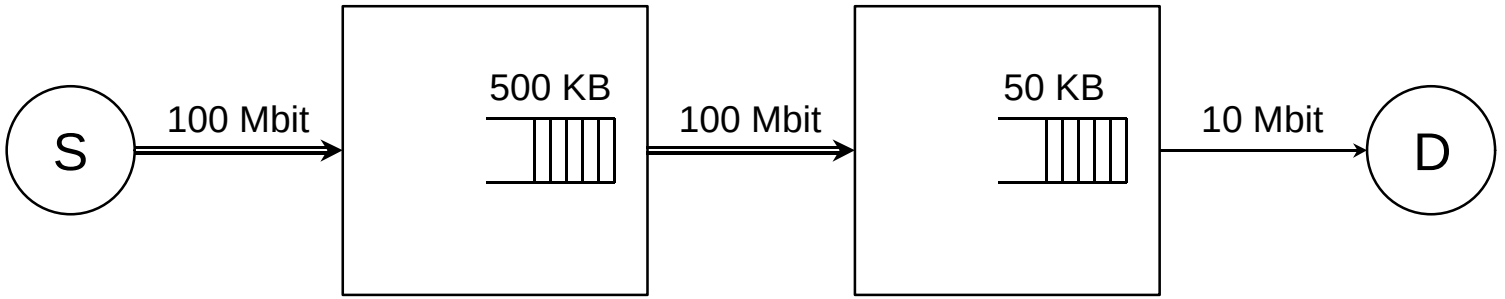

Rate Mismatch

- Traffic
  - Uniform-uniform traffic
    - Burst size range 6,000…96,000 bytes
    - Data rate range 1 Mbps to 10 Mbps
  - Self-similar traffic
    - Burst size range 6,000…48,000 bytes
    - Data rate range 2 Mbps to 8 Mbps
- Congestion control mechanism
  - watermark-based congestion detection
    - low threshold = 70%
    - high threshold = 80%
  - notification information
    - block/unblock with no specific information
- Measures
  - packet loss rate

Rate Mismatch

Packet Loss without Congestion Control

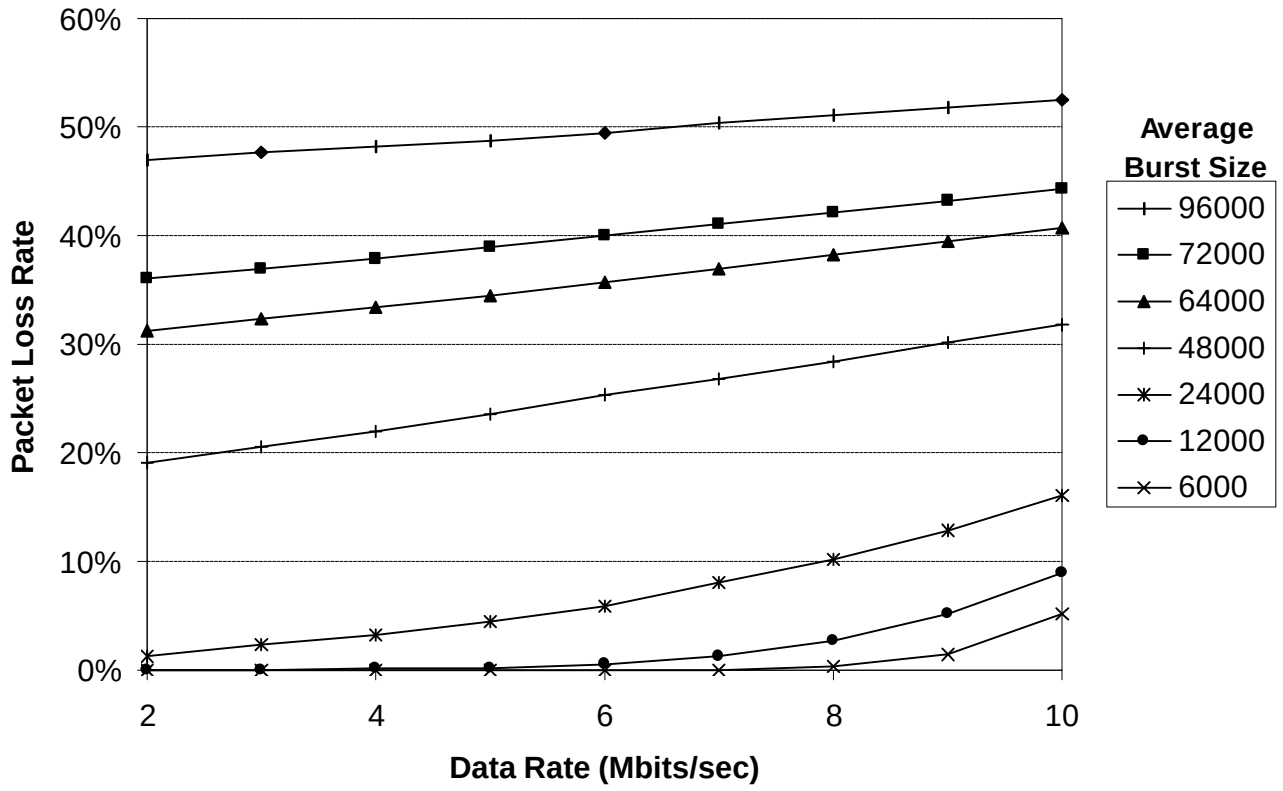

- Uniform-uniform traffic model
- 50 KB buffer becomes congested and drops packets

Rate Mismatch

Packet Loss using XON/XOFF

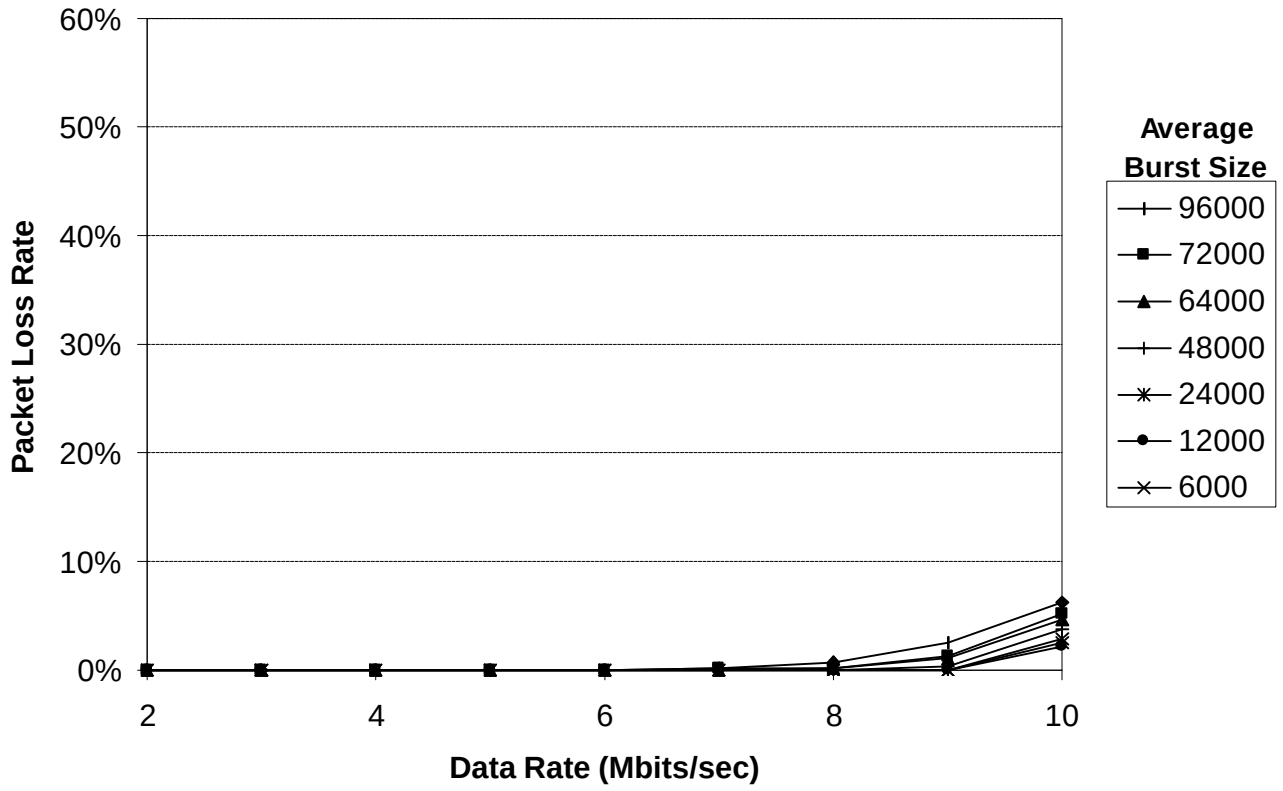

- Uniform-uniform traffic model
- High threshold 80%
- Low threshold 70%

Rate Mismatch

Packet Loss without Congestion Control

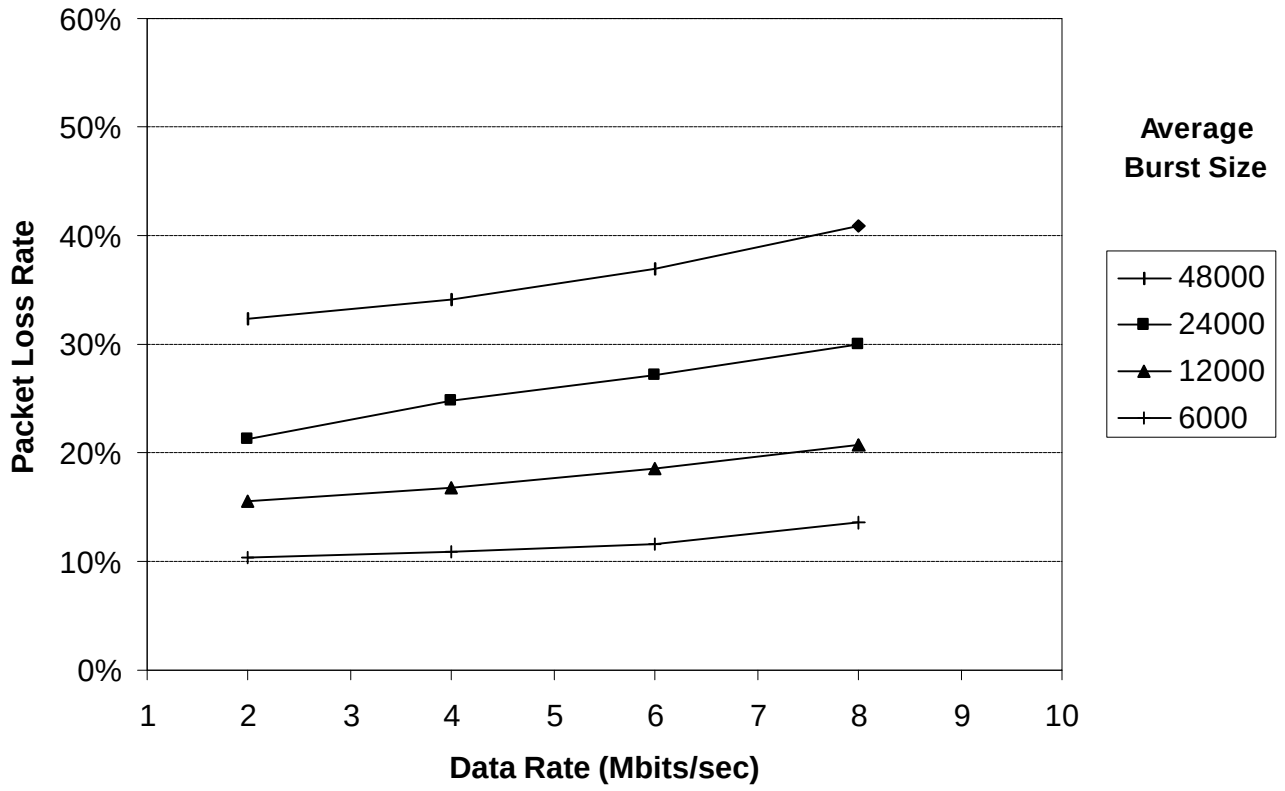

- Self-similar traffic model
- 50 KB buffer becomes congested and drops packets

Rate Mismatch

Packet Loss using XON/XOFF

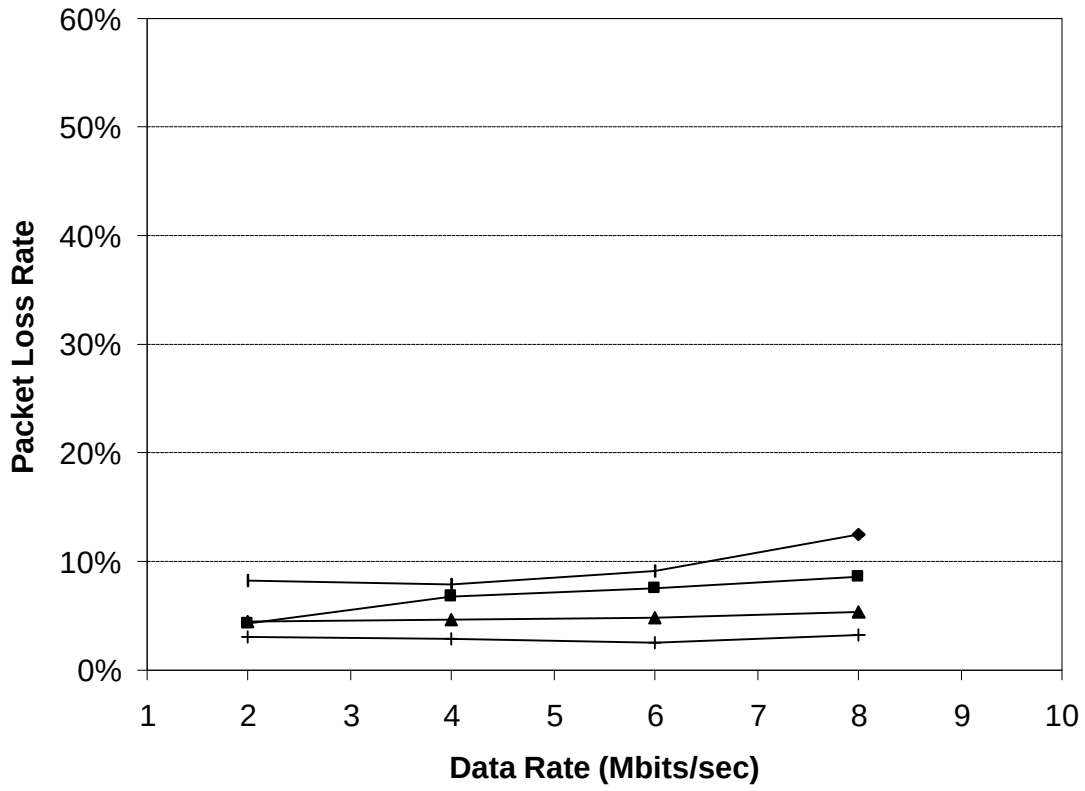

- Self-similar traffic model
- High threshold 80%
- Low threshold 70%
- Can have some loss in 500 KB buffers due to very large bursts

Notes

- Short term congestion may occur in LANs due to bursty traffic
- Congestion control helps reduce packet loss by more efficiently using the distributed buffering resources available in the network

Experiments

Illustrate the need for congestion control based on CoS

B.1 Video and Bursty Data

Video and Bursty Data Topology

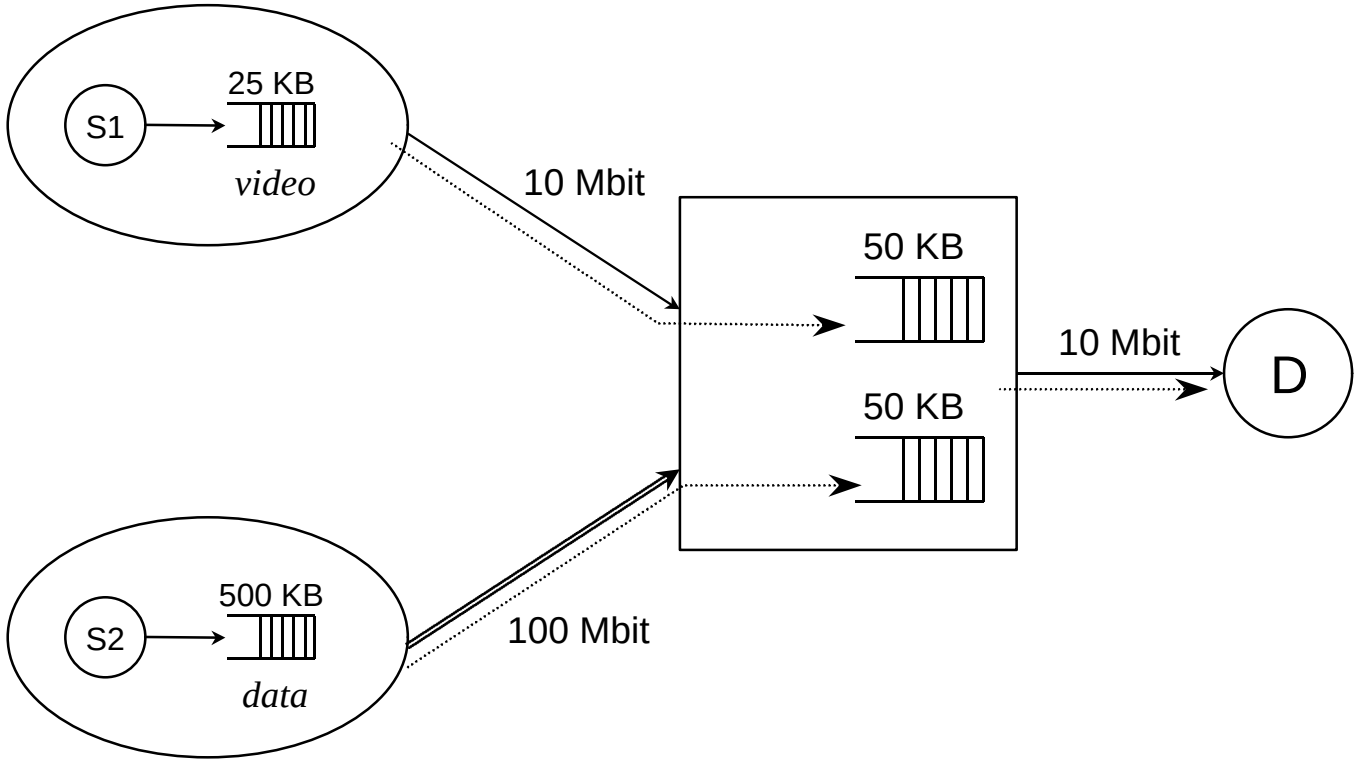

Video and Bursty Data

- Traffic
  - Uniform-uniform traffic
    - Burst size range 6,000…96,000 bytes
    - Data rate range 1 Mbps to 7 Mbps
  - Video
    - 2 streams of 1.5Mbps VBR video
- Congestion control mechanism
  - watermark-based congestion detection
    - low threshold = 70%
    - high threshold = 80%
  - notification information
    - block/unblock with no specific information
- Measures
  - packet loss rate

Video and Bursty Data

Data Packet Loss without Congestion Control

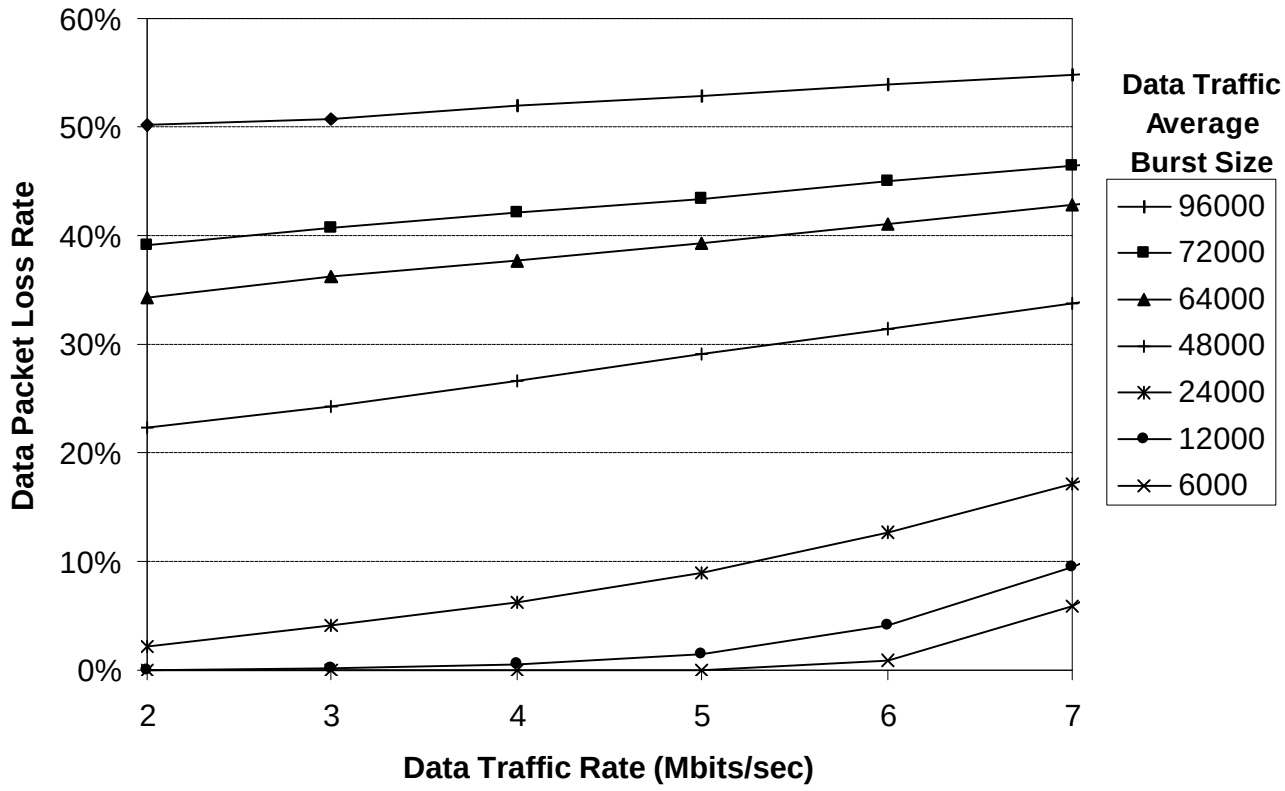

- No congestion control
- Losses occur when 50 KB buffer overflows

Video and Bursty Data

Video Packet Loss without Congestion Control

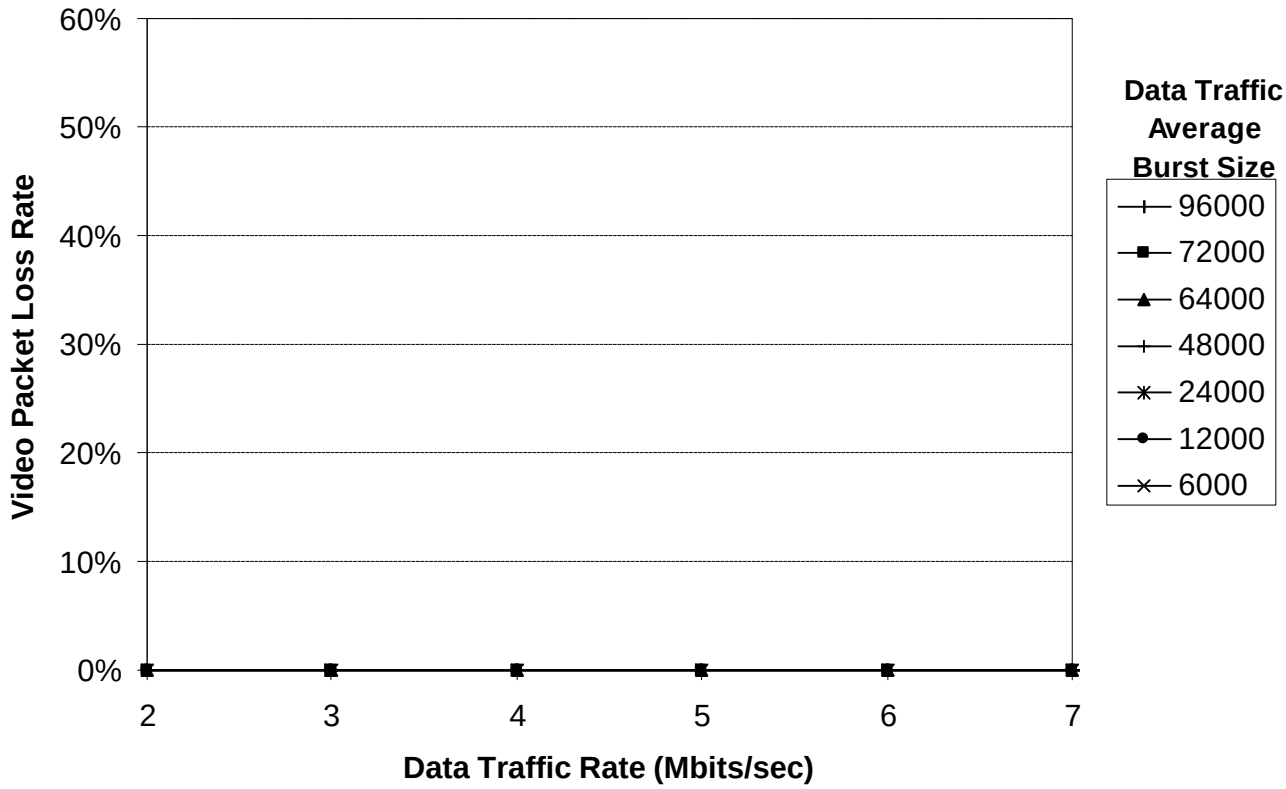

- No congestion control
- Losses occur when 50 KB buffer overflows

Video and Bursty Data

Data Packet Loss using Xon/Xoff

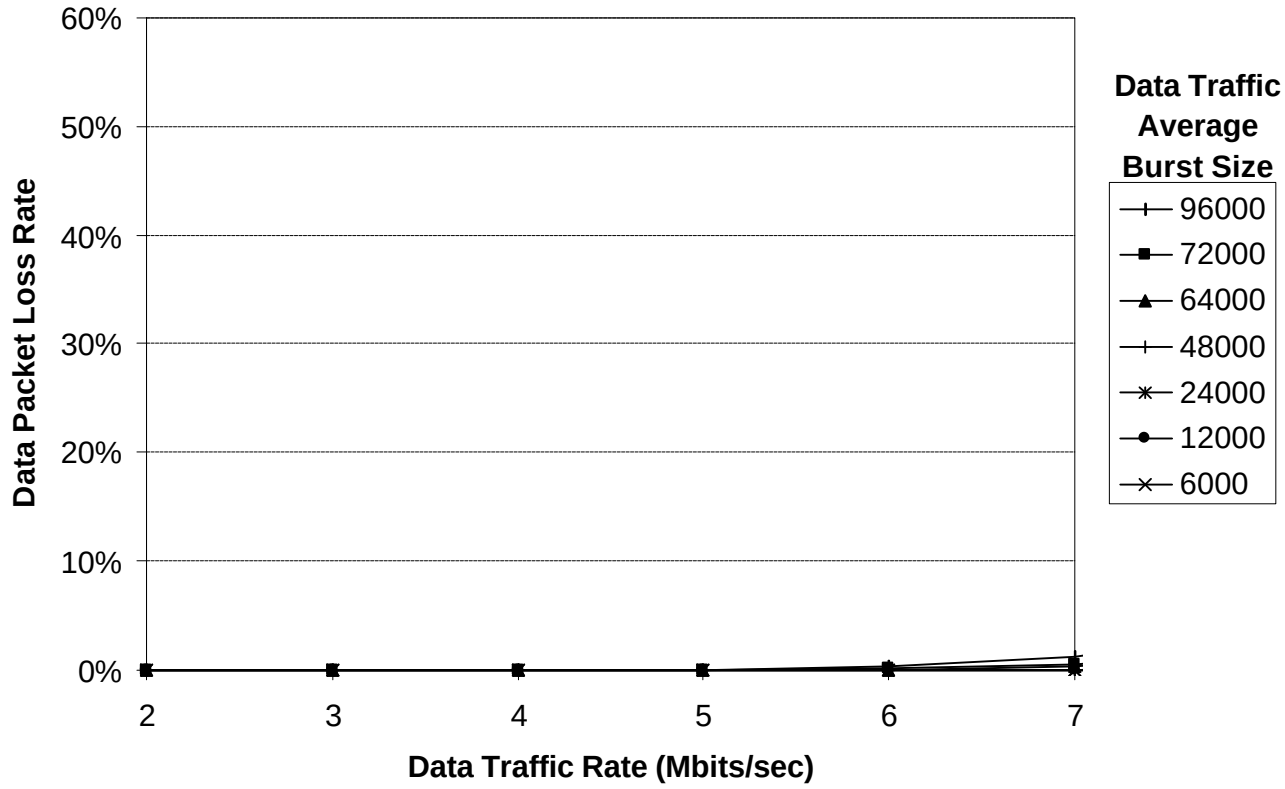

- High threshold 80%
- Low threshold 70%

Video and Bursty Data

Video Packet Loss Using Xon/Xoff

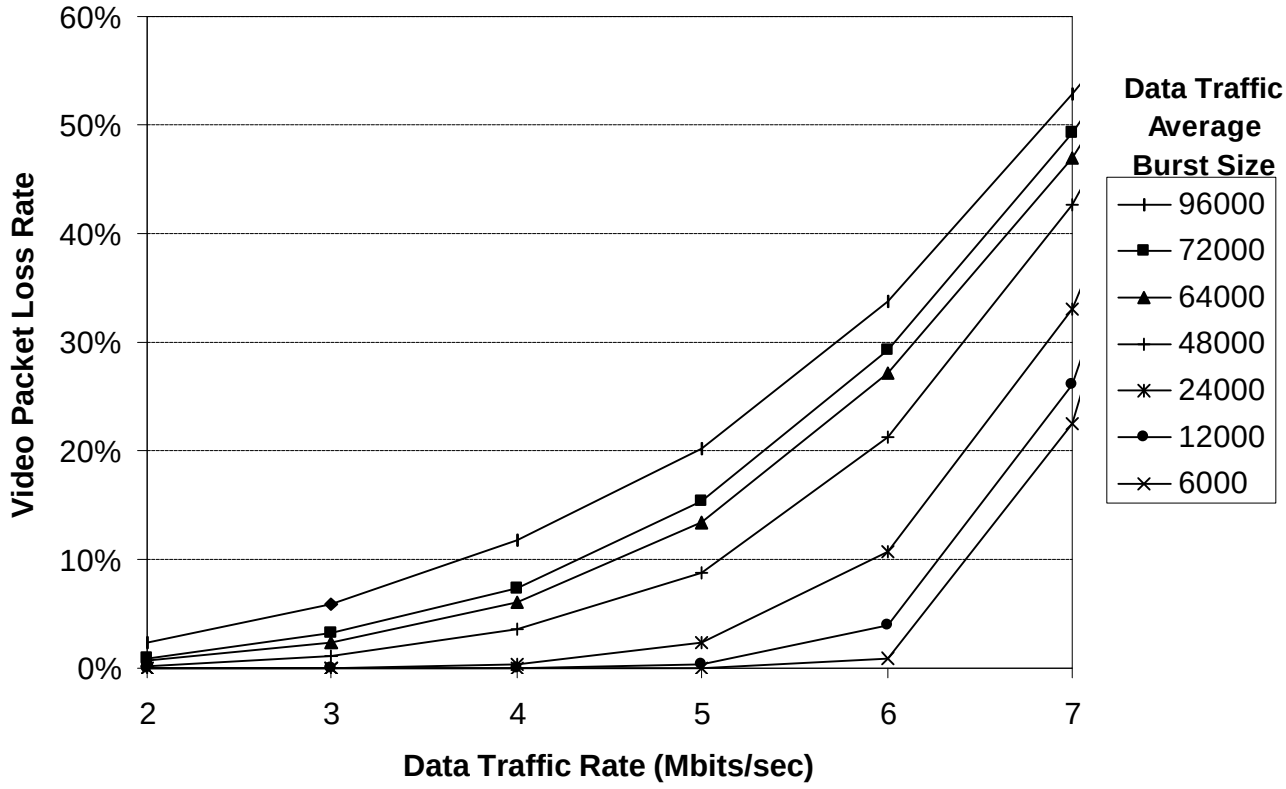

- High threshold 80%
- Low threshold 70%

Video and Bursty Data

Data Packet Loss using Xon/Xoff with Class

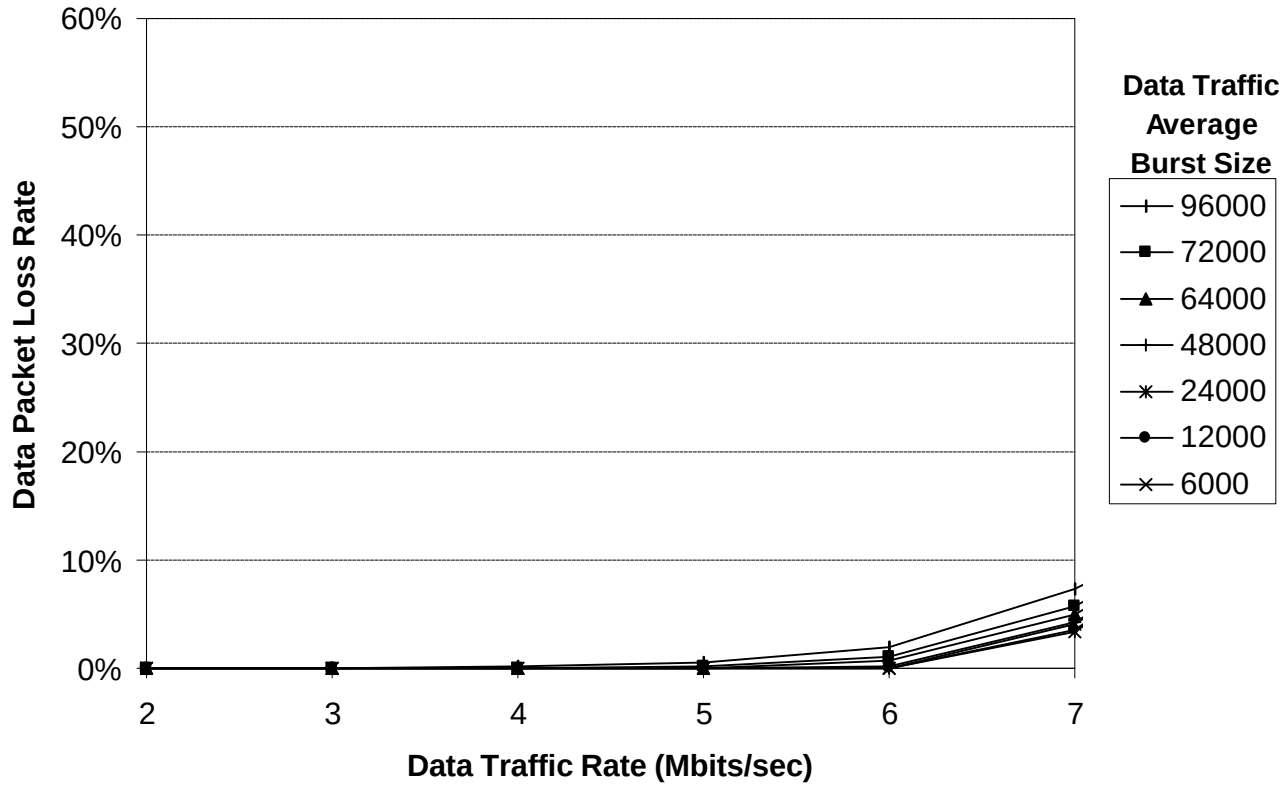

- High threshold 80%
- Low threshold 70%

Video and Bursty Data

Video Packet Loss Using Xon/Xoff with Class

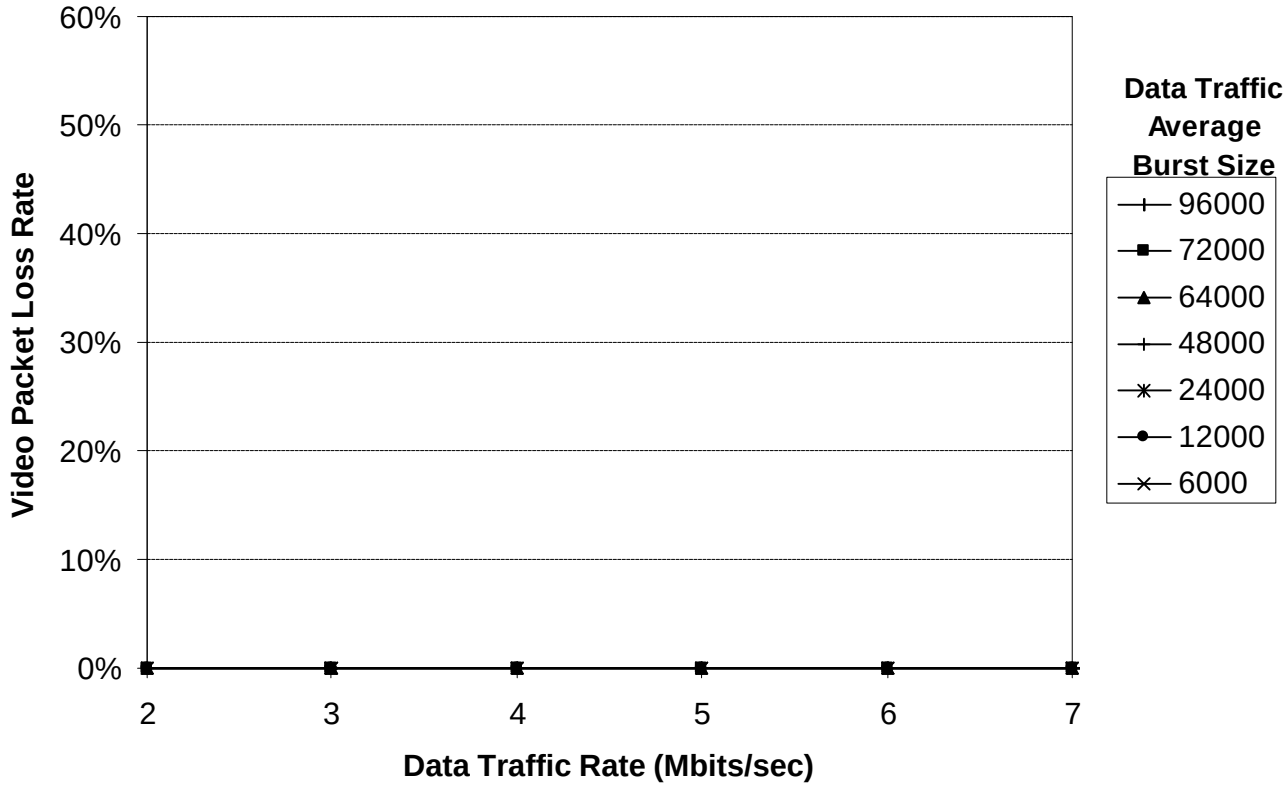

- High threshold 80%
- Low threshold 70%

Notes

- Congestion in a low priority class may severely affect high priority traffic.
- Performing congestion control based on class of service eliminates this problem

Experiments

Illustrate the need for congestion control based on destination address

C.1 Independent flows

C.2 Merging over the backbone

C.3 Source control

Independent Flows Topology

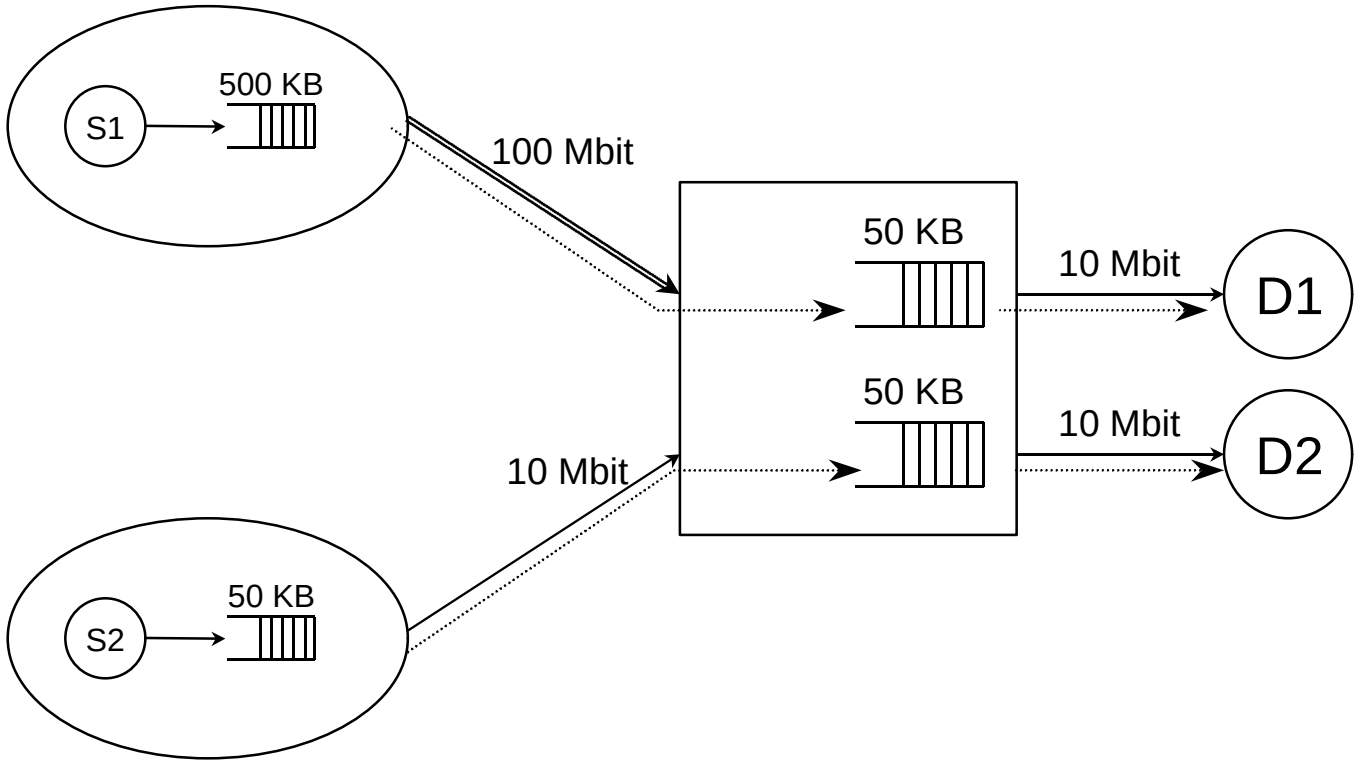

Independent Flows

- Traffic
  - Uniform-uniform traffic
    - Burst size range 6,000…96,000 bytes
    - Data rate range 2 Mbps to 10 Mbps
  - S2 always uses 10Mbps data rate with 6000 byte bursts
- Congestion control mechanism
  - watermark-based congestion detection
    - low threshold = 70%
    - high threshold = 80%
  - notification information
    - block/unblock with no specific information
- Measures
  - packet loss rate

Independent Flows

S1 Packet Loss without Congestion Control

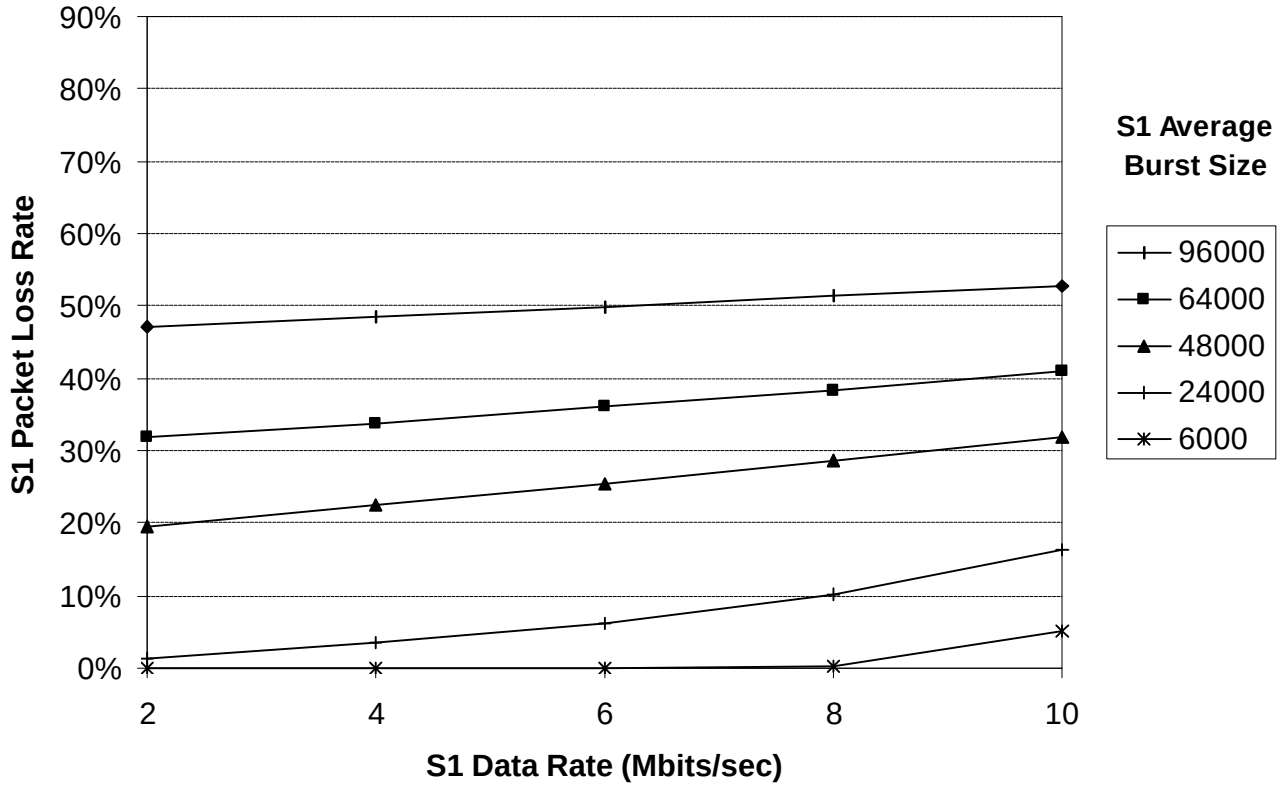

Independent Flows

S2 Packet Loss without Congestion Control

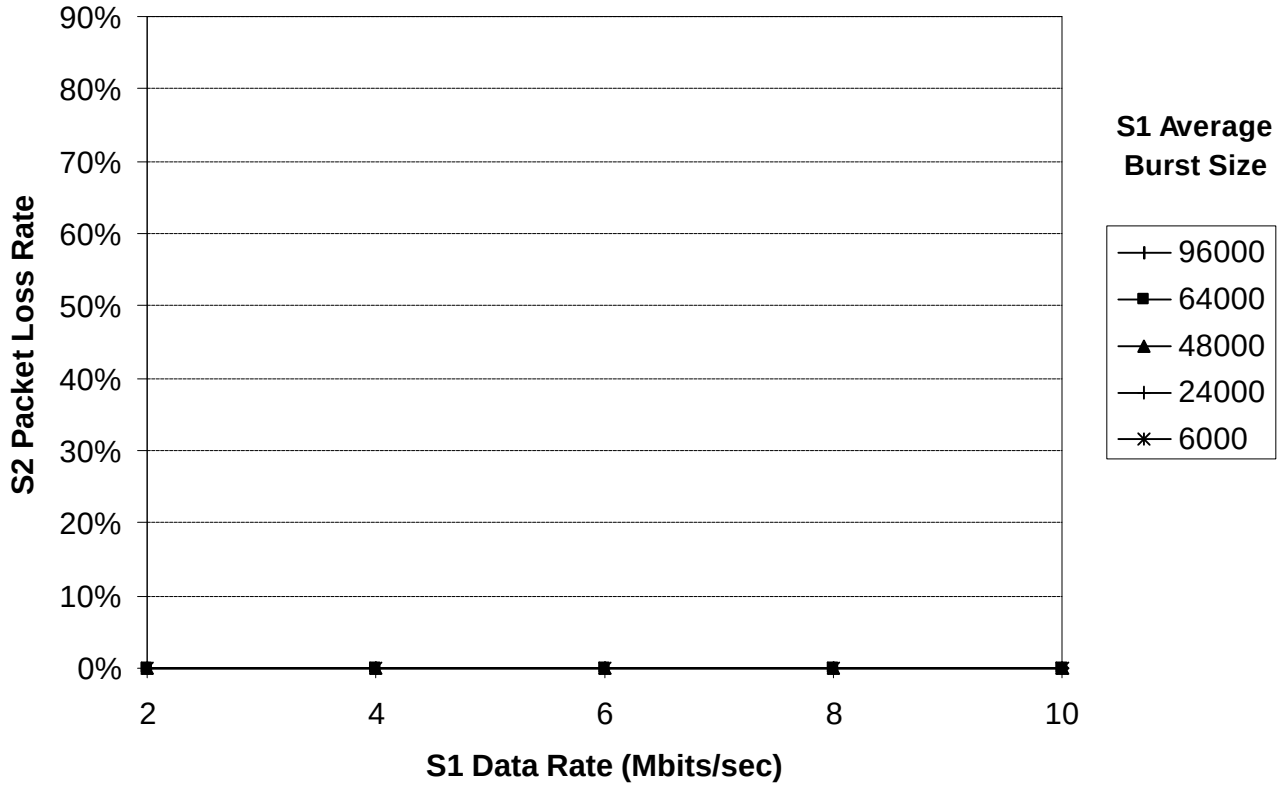

- Plotted vs. S1 data rate

Independent Flows

S1 Packet Loss with Xon/Xoff

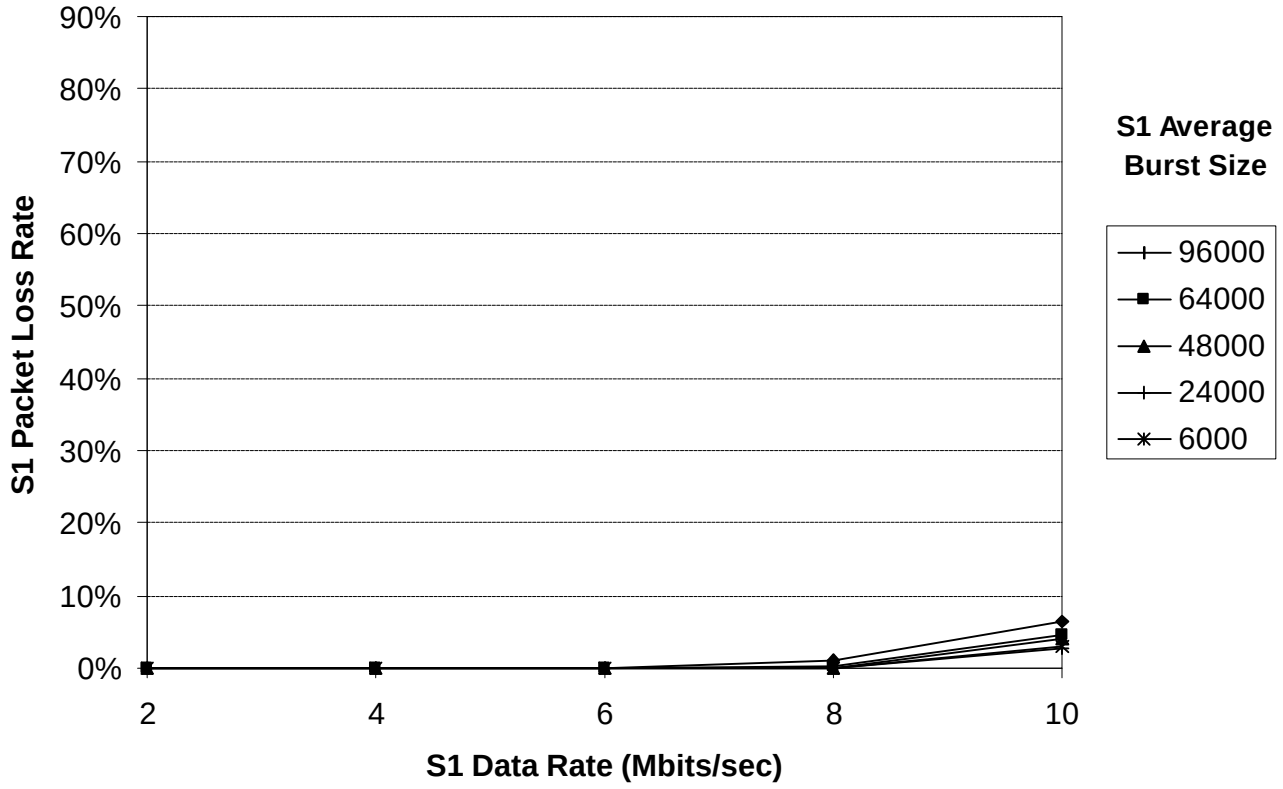

- High watermark = 80%
- Low watermark = 70%

Independent Flows

S2 Packet Loss with Xon/Xoff

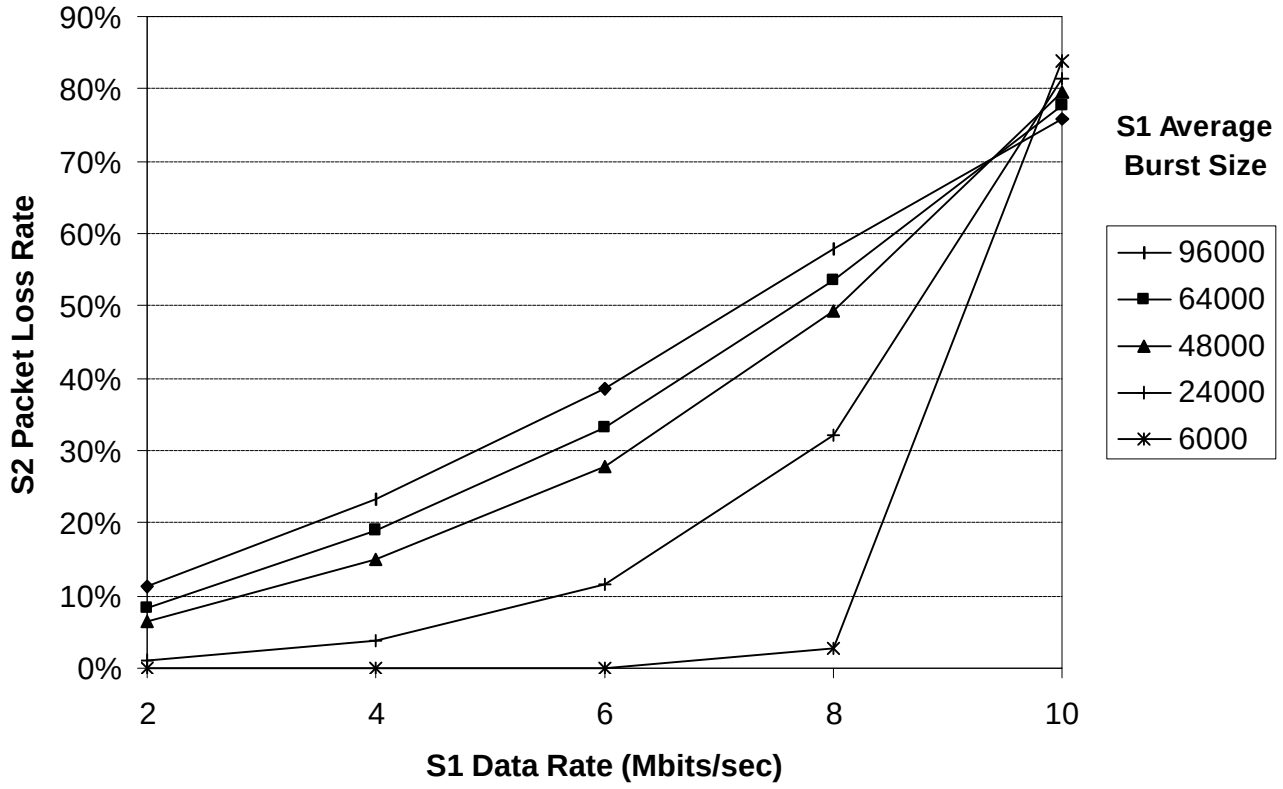

- Plotted vs. S1 data rate
- High watermark = 80%
- Low watermark = 70%

Independent Flows

S1 Packet Loss using Xon/Xoff with Destination

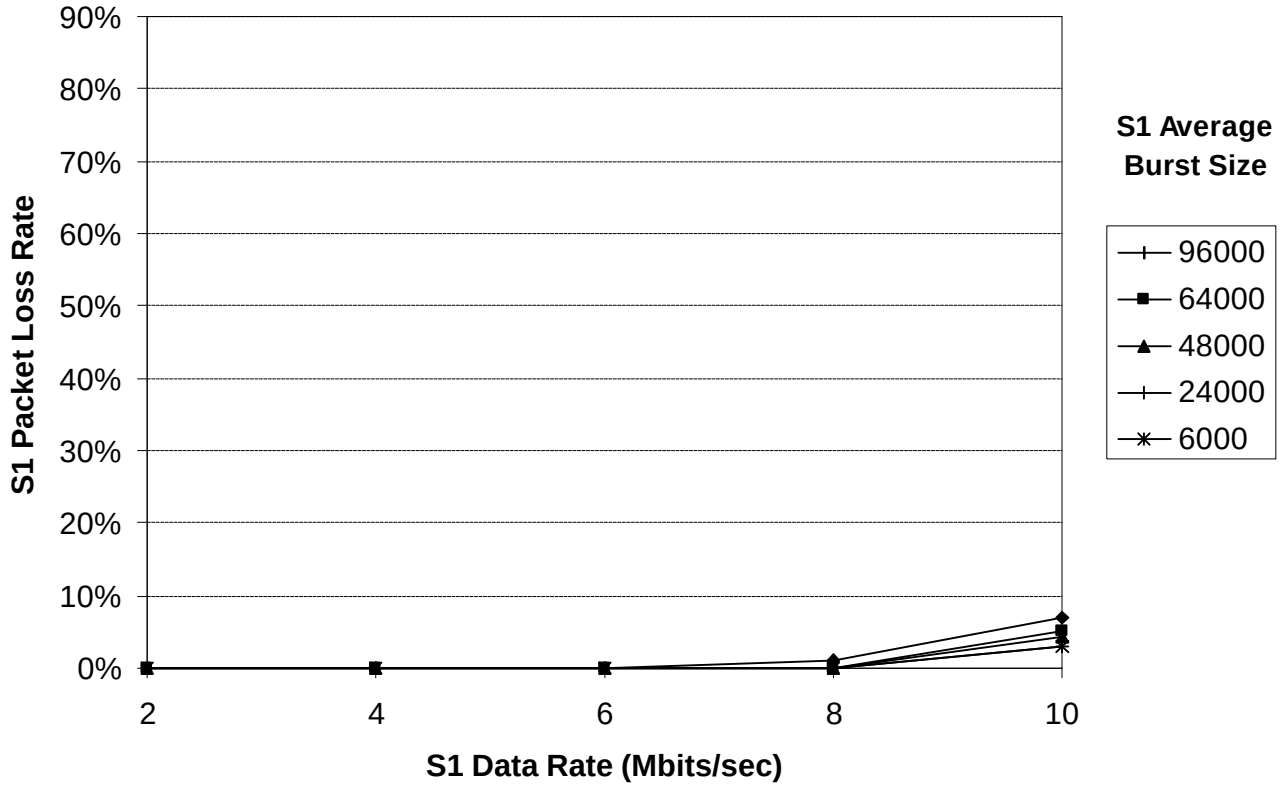

- High watermark = 80%
- Low watermark = 70%

Independent Flows

S2 Packet Loss using Xon/Xoff with Destination

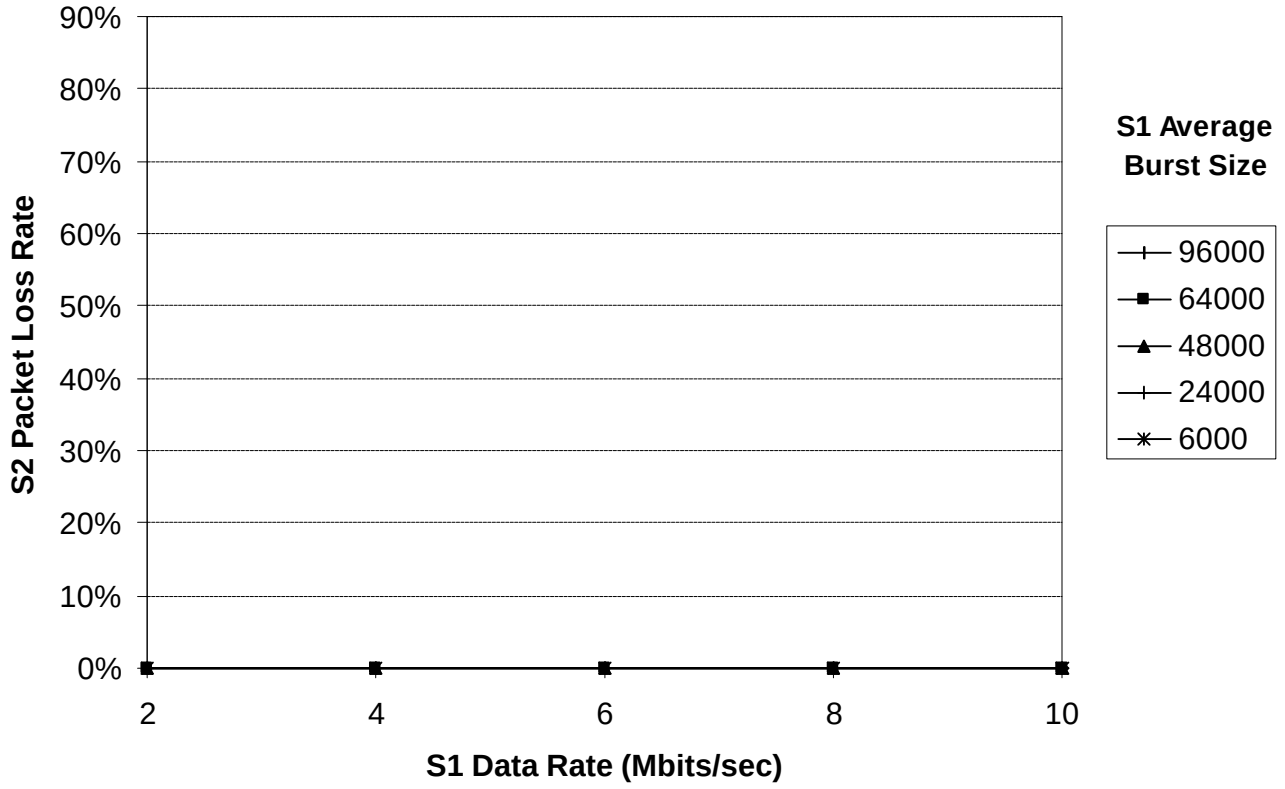

- Plotted vs. S1 data rate
- High watermark = 80%
- Low watermark = 70%

Merging Over the Backbone

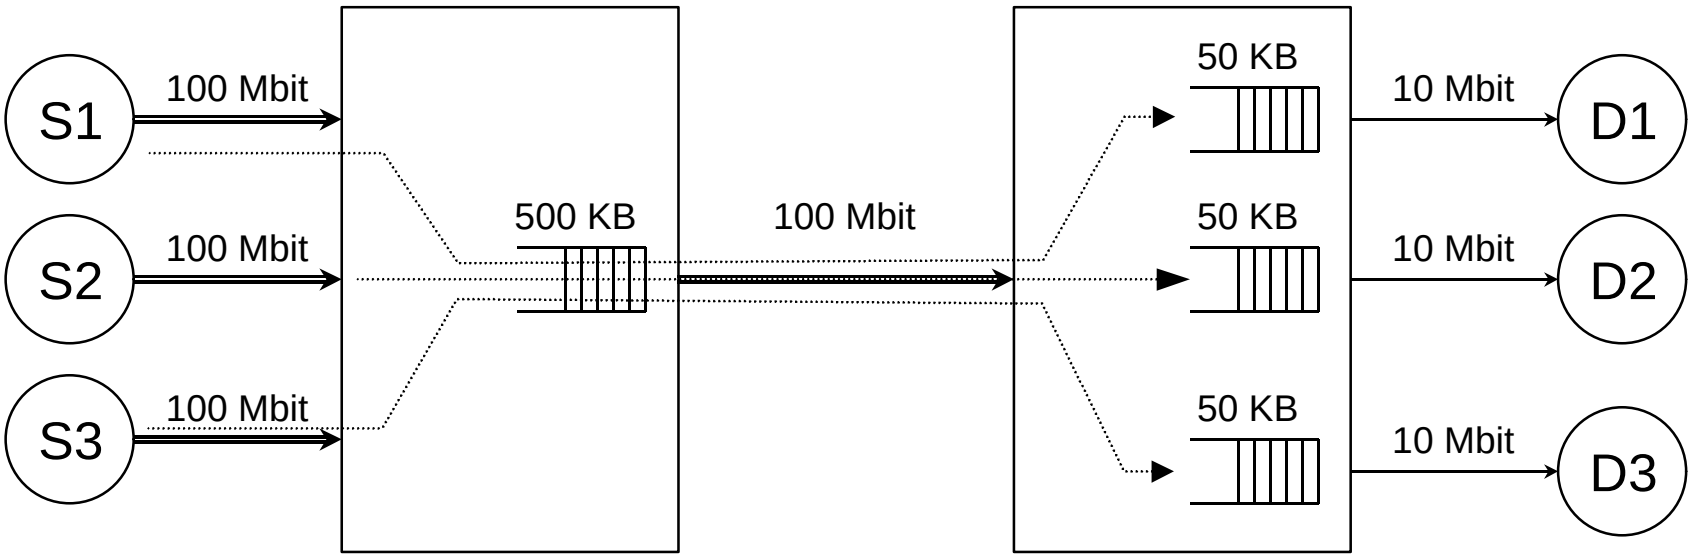

Merging Over the Backbone

- Traffic
  - Uniform-uniform traffic
    - Burst size range 6,000…96,000 bytes
    - Data rate range 2 Mbps to 10 Mbps
- Congestion control mechanism
  - watermark-based congestion detection
    - low threshold = 70%
    - high threshold = 80%
  - notification information
    - block/unblock with no specific information
- Measures
  - packet loss rate

Merging over the Backbone

Packet Loss without Congestion Control

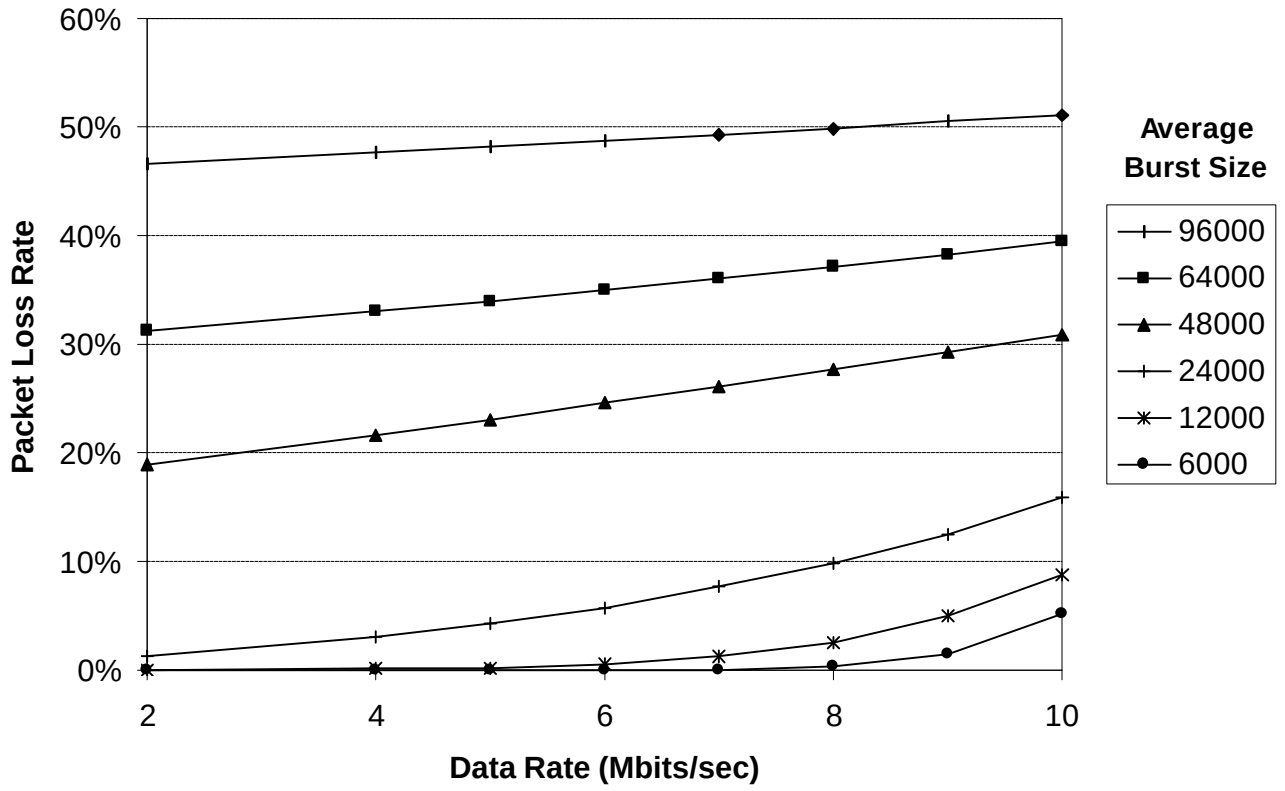

- Uniform-uniform traffic model
- Data rate is given per source
- 50 KB buffer overflows and drops packets

Merging over the Backbone

Packet Loss Using Xon/Xoff

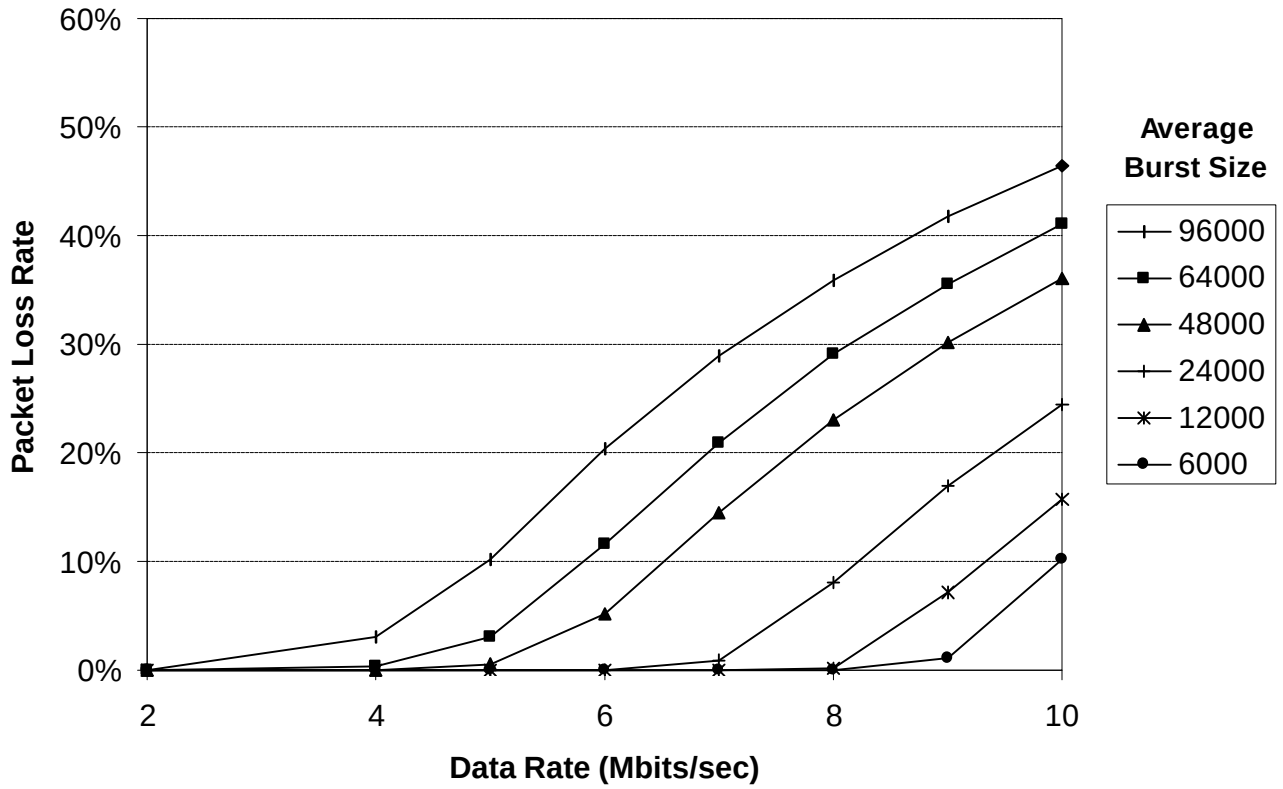

- Uniform-uniform traffic model
- Data rate is given per source
- 500 KB buffer overflows and drops packets at high data rates

Merging over the Backbone

Packet Loss using Xon/Xoff with Destination

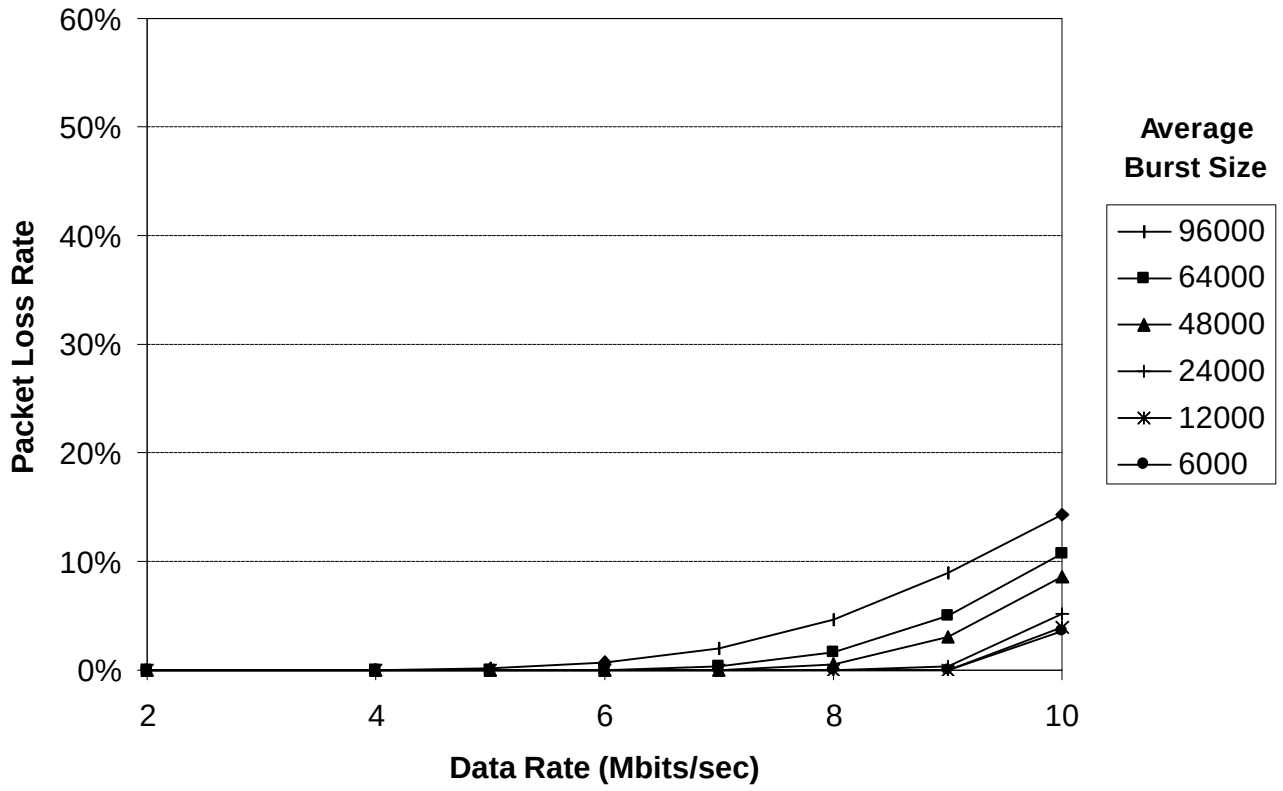

- Uniform-uniform traffic model
- Data rate is given per source

Merging over the Backbone

Packet Loss without Congestion Control

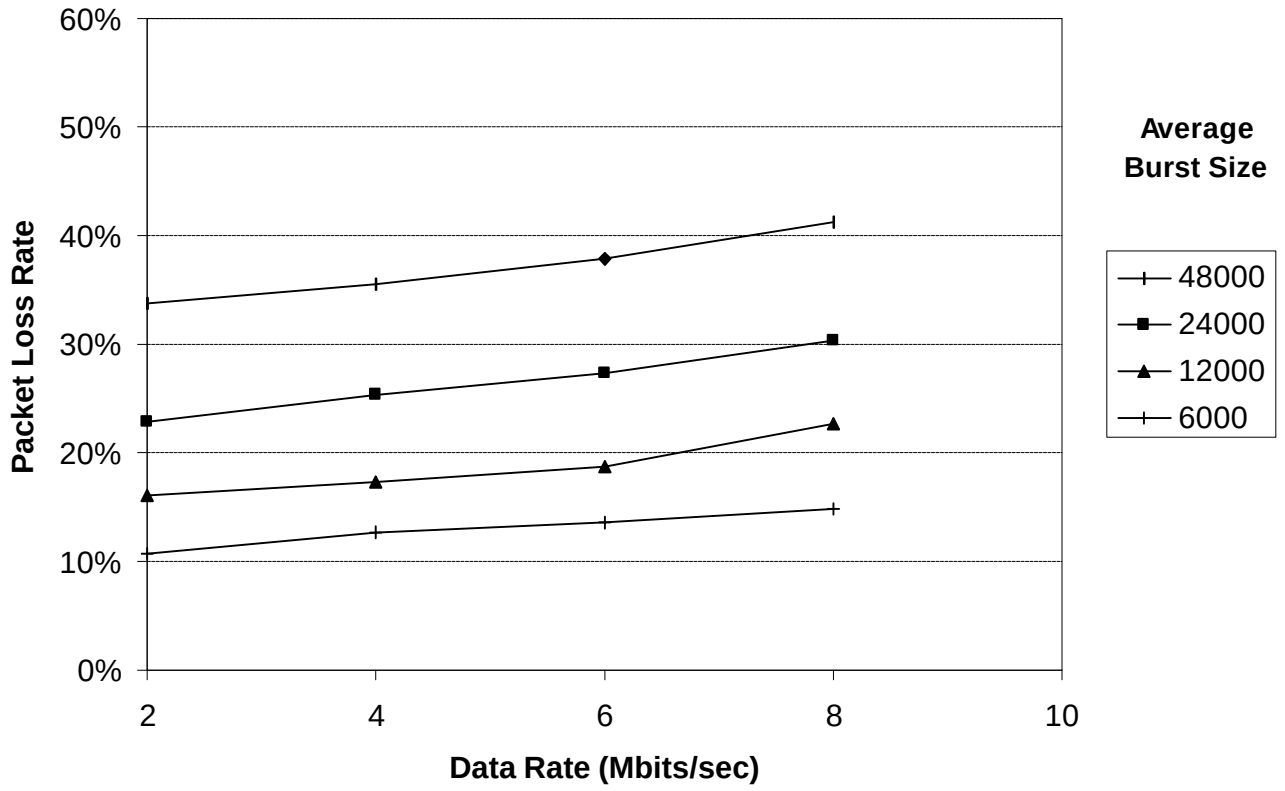

- Self-similar traffic model
- Data rate is given per source
- 50 KB buffer overflows and drops packets

Merging over the Backbone

Packet Loss Using Xon/Xoff

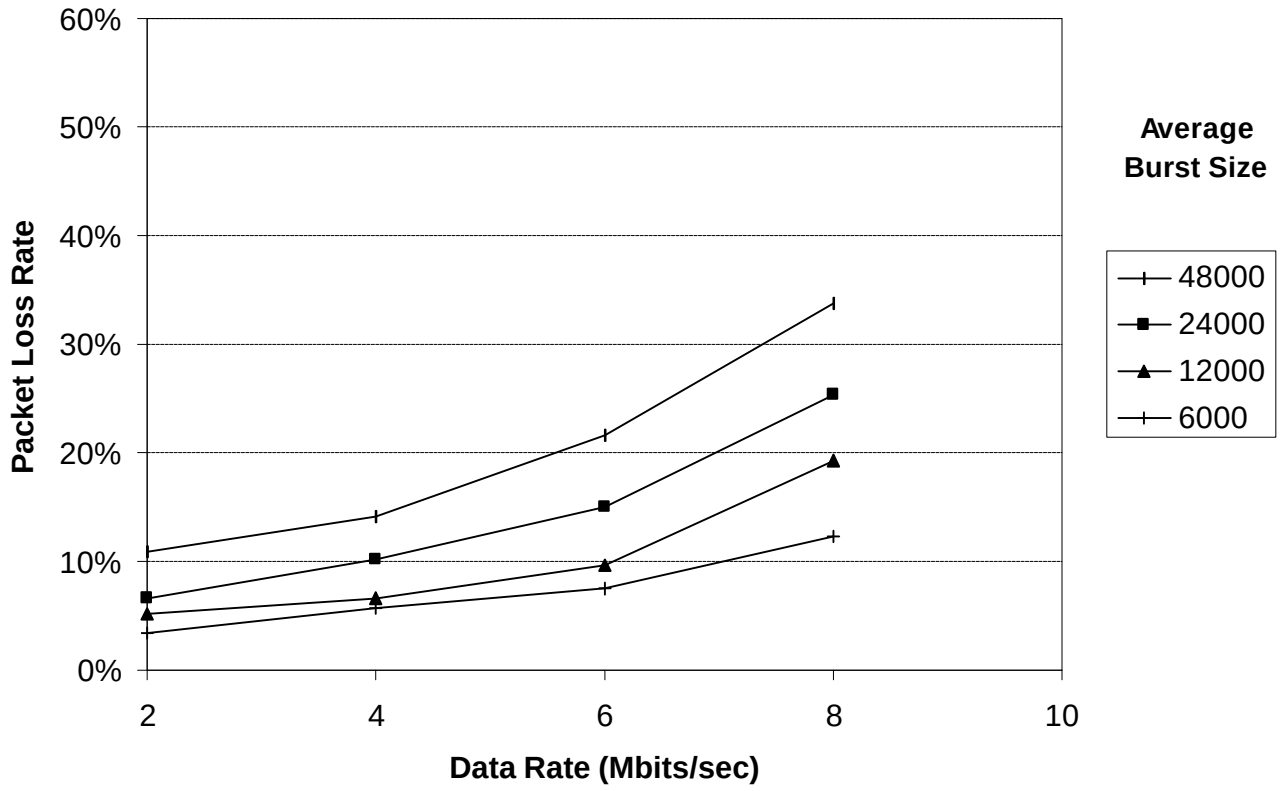

- Self similar traffic model
- Data rate is given per source

Merging over the Backbone

Packet Loss using Xon/Xoff with Destination

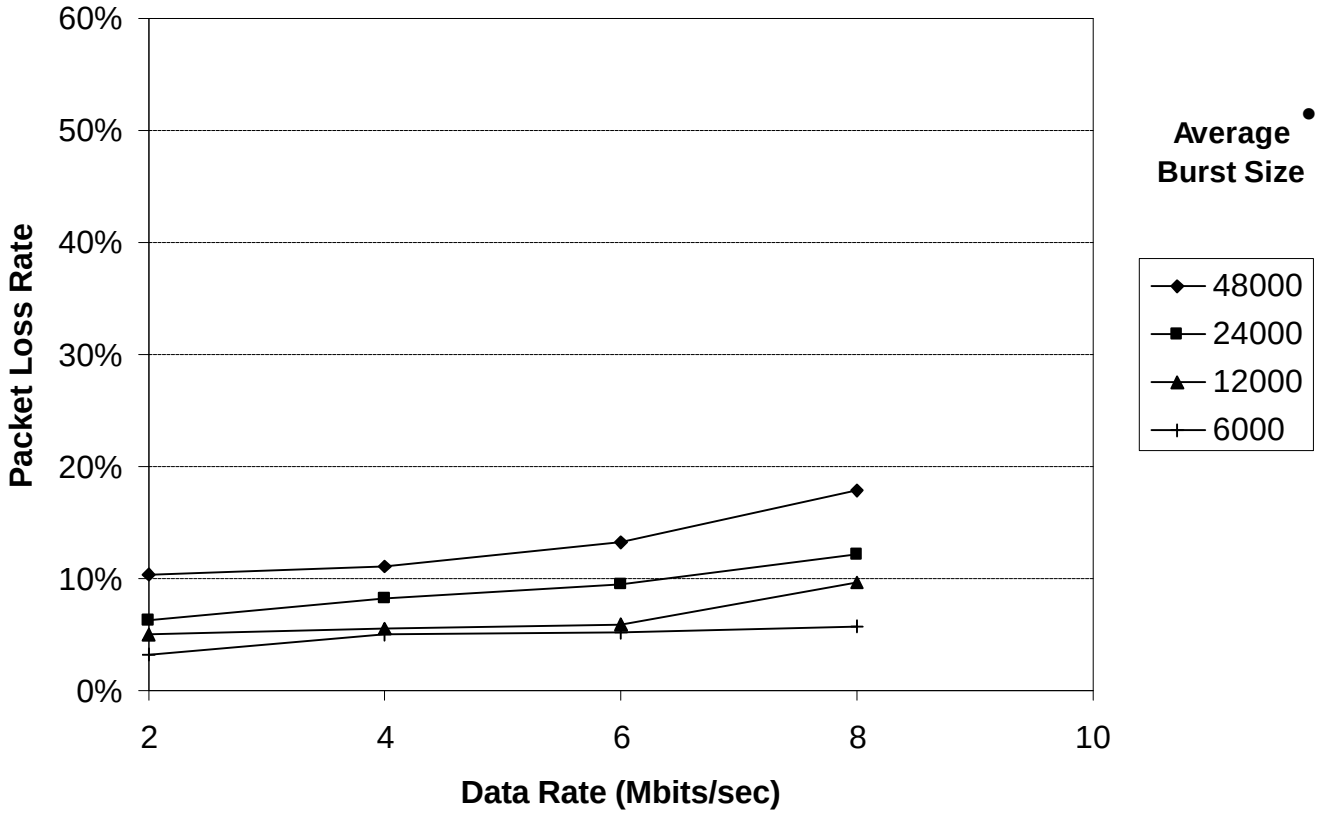

- Self-similar traffic model
- Data rate is given per source

Controlling the Source

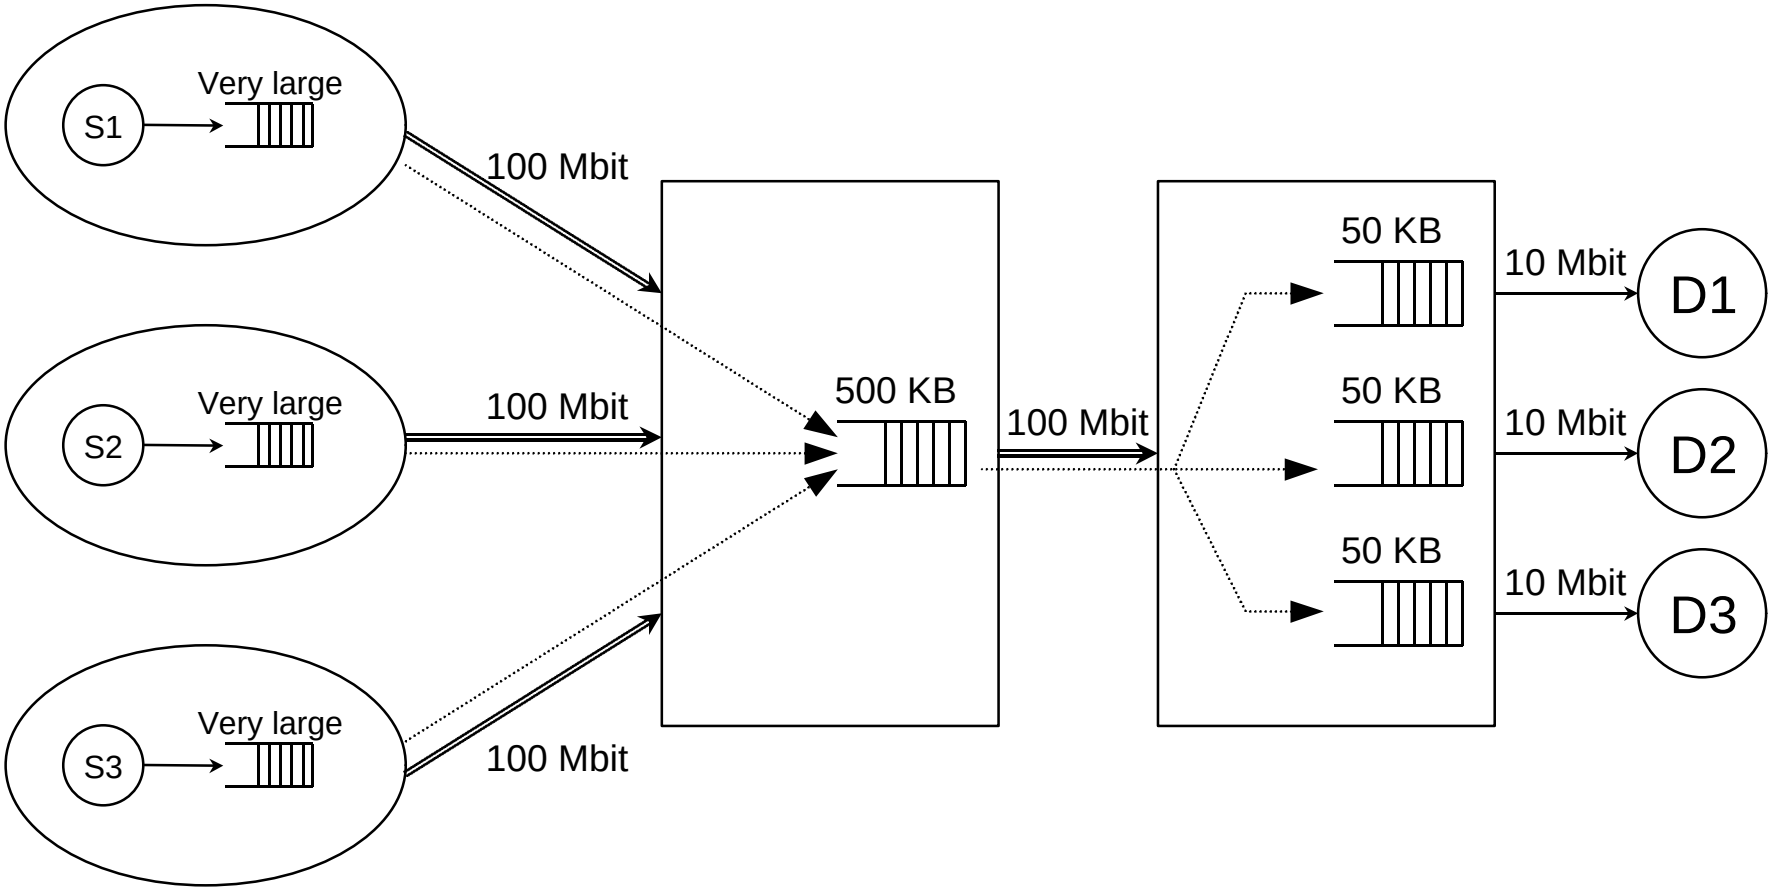

Controlling the Source

- Traffic
  - Uniform-uniform traffic
    - Burst size range 6,000…96,000 bytes
    - Data rate range 2 Mbps to 10 Mbps
- Congestion control mechanism
  - watermark-based congestion detection
    - low threshold = 70%
    - high threshold = 80%
  - notification information
    - block/unblock with no specific information
- Measures
  - packet loss rate
  - delay

Controlling the Source

Packet Loss without Congestion Control

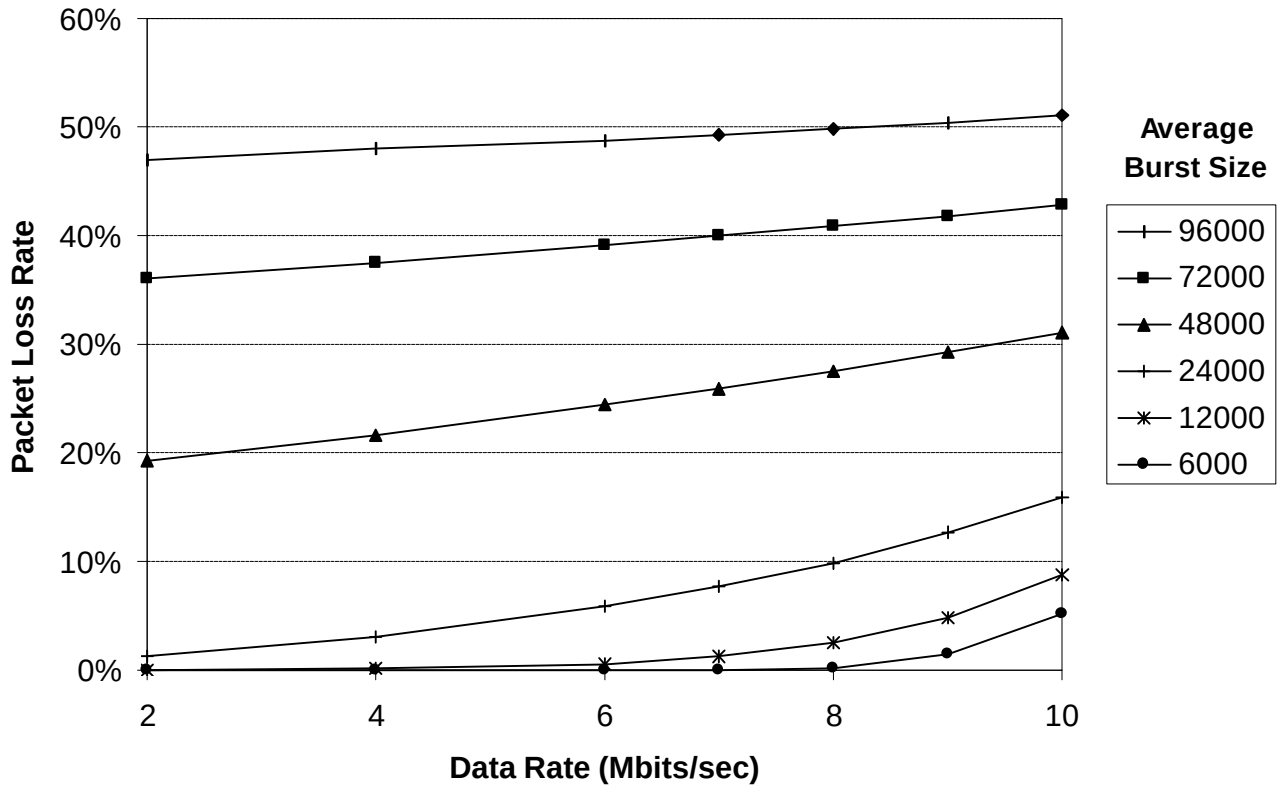

- Uniform-uniform traffic model
- Data rate is given per source
- 50 KB buffer overflows and drops packets

Controlling the Source

Packet Loss Using Xon/Xoff

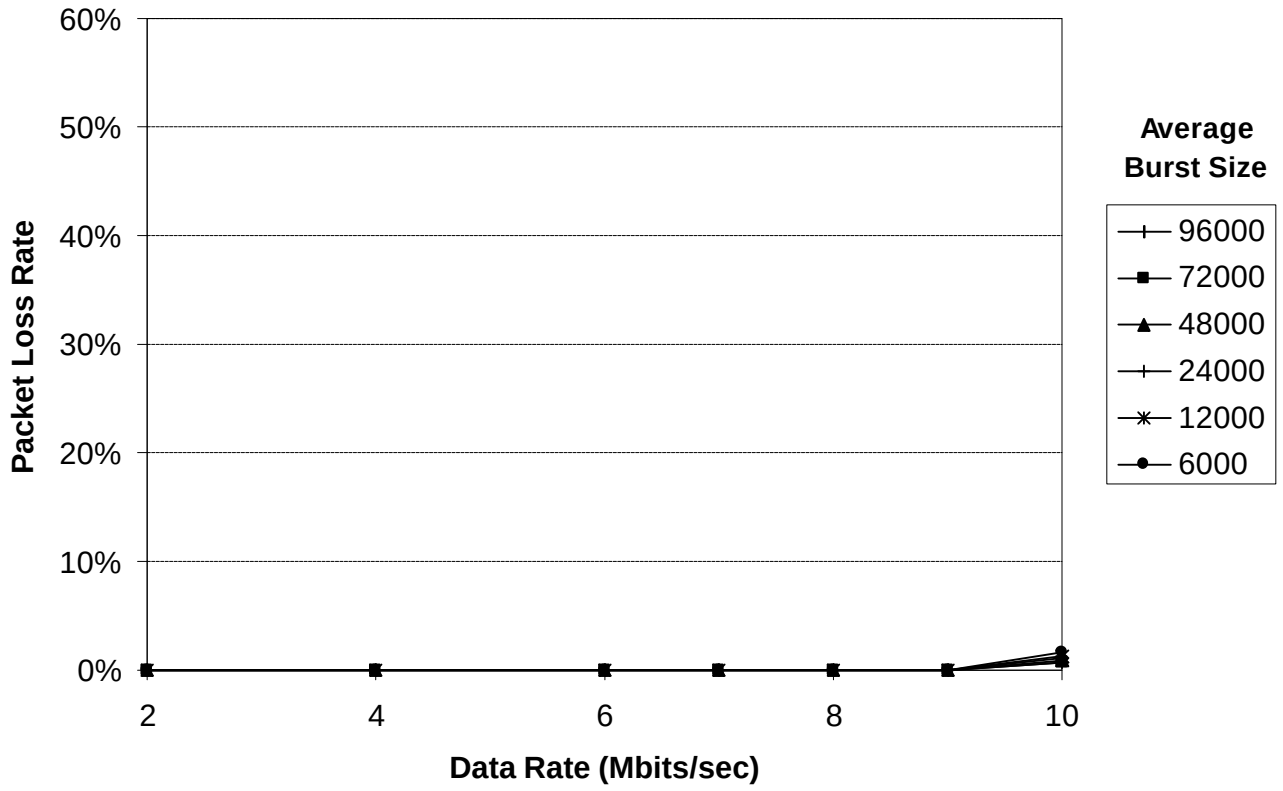

- Uniform-uniform traffic model
- Data rate is given per source
- Extra buffers eliminate loss

Controlling the Source

Packet Loss using Xon/Xoff with Destination

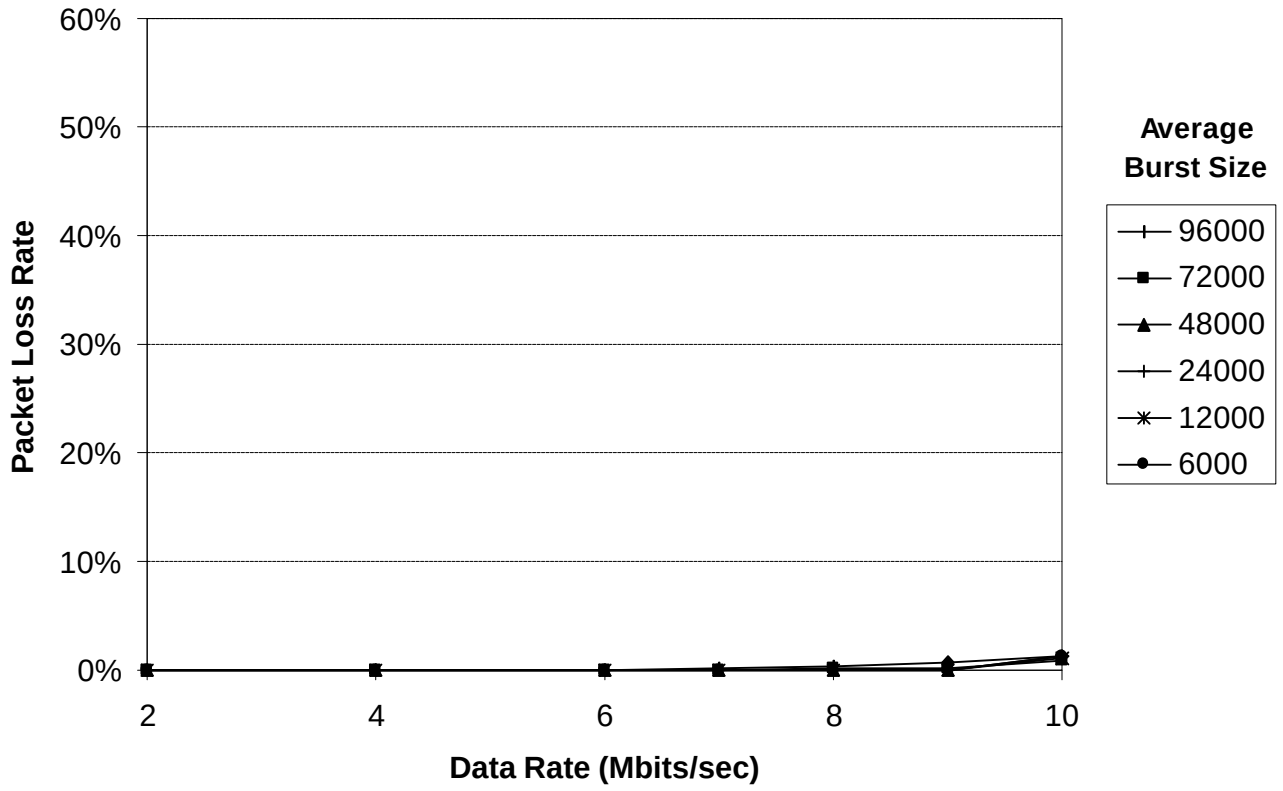

- Uniform-uniform traffic model
- Data rate is given per source
- Packet loss is eliminated

Controlling the Source

Packet Delay

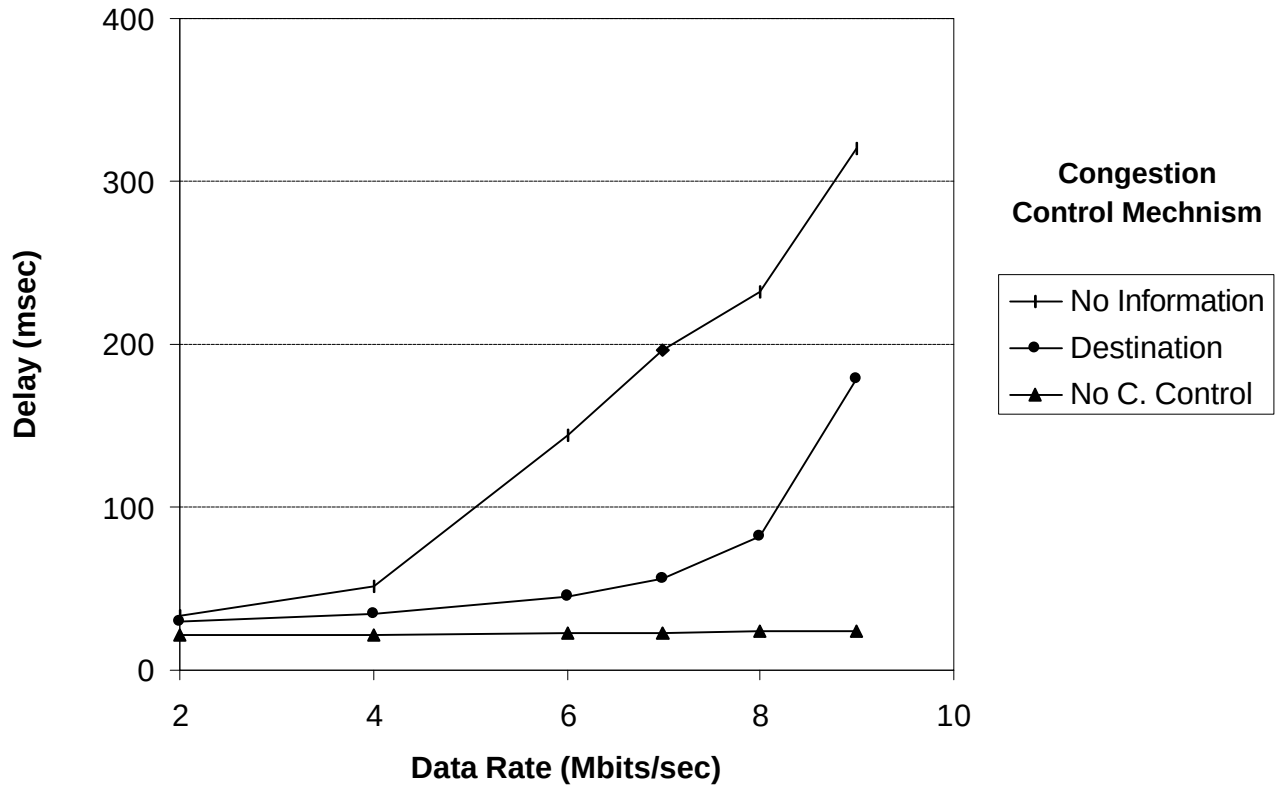

- Uniform-uniform traffic model
- 48,000 byte burst size
- Congestion control based on destination address eliminates loss for a smaller increase in delay

Conclusions

- Congestion detection is necessary
  - Reduces loss due to traffic merging and rate mismatches
- Congestion detection mechanism should include CoS information
  - Without CoS information, congestion of low priority traffic can severely affect high priority traffic
- Congestion detection mechanism should include destination address information
  - Not using destination information can limit the achievable throughput of the network and increase packet delays